An Architecturally Relevant Model for Creating Orientation Maps of Primary Visual Cortex

BY

JENNIFER E. ANDERSON
M.A. Psychology (University of Illinois at Chicago) 2010
B.A. Psychology (University of Illinois at Chicago) 2007

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois at Chicago, 2013

Chicago, Illinois

Defense Committee:

Tanya Berger-Worl, Chair and Advisor
Robert Kenyon
Barbara DiEugenio

I dedicate this work to all women in Computer Science.

ACKNOWLEDGMENTS

I would like to thank my thesis committee – Dr. Tanya Berger-Wolf, Dr. Robert Kenyon, and Dr. Barbara DiEugenio – for their support and assistance in completing this project. I would also like to thank both the Department of Computer Science and Department of Psychology, Division of Behavioral Neuroscience, for the opportunity to pursue both disciplines concurrently during my graduate career. Without these opportunities, I would not be able to follow my dream of working on research questions within the niche domain of Computational Neuroscience.

A number of individuals were extremely helpful to me with their ideas and insight, and I would like to thank them as well – Anushka Anand, Researcher at Tableau Software; Simone Franzini, PhD, Software & Analytics Technology Developer at Motorola Mobility; Michael Levine, PhD, Professor Emeritus of Psychology (and my Advisor in Behavioral Neuroscience) – all of whom I have met during my graduate career at UIC.

JEA

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

dLGN	Dorsal Lateral Geniculate Nucleus
MGM	Medial Geniculate Nucleus
V1	Primary Visual Cortex
SOM	Self Organizing Map
CW	Clockwise
CCW	Counter-clockwise
MST	Minimum Spanning Tree
CV	Coefficient of Variation
PCA	Principle Component Analysis
LTP	Long Term Potentiation

SUMMARY

Building biologically relevant models is a area of research experiencing a boom in growth. Specifically, many companies and research teams are trying to understand a very efficient and powerful organ of the body – the brain. The brain weighs 3 pounds consuming a mere 20 watts of energy, yet it is able to perform seemingly thousands of tasks, whether they be visual, cognitive, or hormonally driven, all simultaneously with the utmost efficiency.

Understanding how the brain works is an important milestone for humankind, but it is far from easy. One starting point is to try and model the input-output responses of the brain and somehow formalize these response behaviors. When creating computational models, it is important to try and stay as true to the biological model as possible. This includes constraining the model to the architecture and the wiring of components of the biological model. The reason we constrain ourselves to the architecture is because we need to allow for any unaccounted for properties to develop within the model. Emergence, for example, is a one of these properties; it is the way complex patterns arise out from some multiple of relatively simple interactions. Emergent properties are common in the brain (intelligence being the most commonly cited).

The visual system comprises nearly one-third of all of the processing within the human brain and is also one of the most widely studied areas of the nervous system, thus it is appropriate to try and create a model from this plethora of data. In this thesis, I demonstrate that by following the basic rules of brain organization – sheets of neurons working in parallel organized into a

SUMMARY (Continued)

specific hierarchy – I can achieve a biologically meaningful output. This model is composed solely of leaky integrate-and-fire neurons constrained to a specific layout.

In the end, I show that this biologically-inspired model can produce outputs equivalent to empirical data. Additionally, the architecture allows for emergent properties to develop (i.e. Long Term Potentiation between neighboring neurons). I detail the parameters of the model and showcase the parameters that can generate biologically similar maps. These comparisons are done across a multiplicity of dimensionless measures.

CHAPTER 1

INTRODUCTION

Models are typically developed as a way to help explain a system. Using models, we can study the effects of different components within a system and can also make predictions about certain behaviors produced by that system.

The brain is often referred to as a biological super-computer. In the interest of understanding how the brain works, various models of the brain have been developed (1). Over one-third of the brain is devoted to processing of the visual world and thus, the visual system is one of the most widely studied and well documented areas of the brain (2).

Despite these advances, the extent of the current knowledge is comprised in basic input-to-output models. That is, neuroscientists are able to study the brain empirically, essentially treating it like a "black box". The neuroscientist feeds the brain system a specific input and simply records a response as the output. Through this way, input-output patterns can be mapped, and behavior predicted (3). However, this input-to-output mapping approach is unable to describe how a particular input is transformed into an output. In order to be useful, a model should do more than simply predict a behavior or response, it should be able to provide insight into the inner workings of the system under study. Many of the current models of prediction are abstract in how they achieve a certain output. There usually is no biological basis of reasoning for their design, and any attempt to relate a model back to biology commonly reads as *ex post facto*.

One often overlooked property of the intelligent brain is that of emergence. Emergence is the way complex systems and patterns arise out of a multiplicity of relatively simple interactions (4). Emergence is central to the theories of integrative levels and of complex systems, yet many models of the brain fail in this respect (4). Current brain models are limited in that they only map an input to an output, with predictions resulting from patterns based on these mappings. However, an emergent property cannot be localized (5) and thus, it cannot be fed an input nor can an output be recorded from it. Models must be built to allow for the development of emergent properties, and to do this, it must be kept simple without making biologically outlandish assumptions. For example, one may find that a particular algorithm creates an output that appears similar to biological data, but if the algorithm has no basis in biology, then the framework for important emergent properties does not exist, and thus the model itself is not biologically plausible.

1.1 Problem Statement

This thesis demonstrates that a biologically plausible output can be generated with a model built upon biological principles. That is, this model will produce output similar to empirical data using two basic architectural properties of the brain: parallel and serial processing. This architecture provides a framework for the development of relevant emergent properties.

In terms of parallel processing, signaling units called “neurons”, which are based off a simple computational model of the biological neuron, are organized into sheets. For serial processing, these sheets of neurons are organized in a specific architecture in accordance with the biological brain. The result demonstrates that even a simple model adhering to only these two principles can create biologically plausible outputs. Importantly, the inner-workings of the model are transparent (not abstract or unknown), and biologically un-related algorithms have not been assumed and/or programmed into the model. All relevant properties emerge solely as a result of the architecture.

CHAPTER 2

BACKGROUND

The brain is an organ of the body weighing approximately three pounds (6). It is a structure responsible for much of our behavior as well as our ability to interpret and to successfully understand the world in which we live. Because of it, we are able to think and make decisions; we interpret a set of percepts and act upon them in the most optimal manner that we know of at that moment.

Given this, the study of the brain and its computational underpinnings has been a topic of much research. Published academic papers within the field of Artificial Intelligence has seen explosive growth, and with the advent of large scale projects such as synthetic brain simulator Blue Brain Project (7), or seemingly intelligent agents like Deep Blue (8), research in this area will only continue.

However, before one can start simulating the brain, one must understand its architecture. The brain is made up of units called neurons. Neurons are the signaling units of the nervous system. Their anatomy allows them to receive incoming binary pulses via their "dendrites". Dendrites are small extensions connected to the neuron's cell body, wherein the cell's DNA and other important processes are located. At the base of the body extends the axon, a myelinated structure which helps propagate the neuron's own signals downstream. The axon hillock, connecting the axon to the cell body, is what might be considered the equivalent of the "step" in a step-function. That is, the axon hillock collects signals from the dendrites, each of which slightly modify the electrical potential of the cell. If the electrical charge of the cell body is increased faster than the cell can work to maintain its resting equilibrium (called its “resting potential”), then the neuron will depolarize, sending an impulse signal of its own through the axon. The axon ends at the neuron’s “terminal buttons” which is where other neurons can connect to receive the current neuron’s pattern of instructions (see Figure 1) (5).

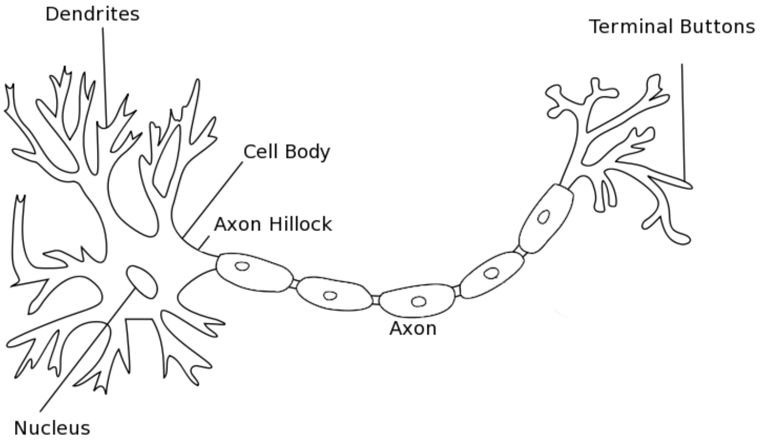

Figure 1. Diagram of a neuron.
Dendrites receive incoming signals and relay them to the cell body. The nucleus contains the neuron’s DNA. The signals are collected at the axon hillock, where the firing threshold is located. Signals are send down the axon to the terminal buttons, where post-synaptic cells can connect.

This is a very basic and assuming description of how a neuron would operate computationally. There are other factors, such as chemical reactions that occur within the synapse at receptors, as well as secondary messengers that work from within the cell; both allowing for delayed integration of a signal into the step-function. Additionally, not all neurons of the nervous system produce action potentials. For example, there are neurons in the retina that are so close together, they do not need to send action potentials to their post-synaptic neighbors (2). Instead, these neurons rely solely on the graded potential changes that occur due to changes in levels of neurotransmitter. These properties allow for a plethora of different patterns to be generated by a single neuron.

Neurons can connect to each other forming tight, interconnected networks. While the basic biology of neurons is the same, how they interconnect with each other dictates how they process a set of signals. Within the brain proper, such groups of neurons are called "nuclei" (outside of the brain proper, such as within the spinal cord, these groups are called "ganglia"), and they perform very specialized tasks. For example, the dorsal Lateral Geniculate Nucleus (dLGN), located within the hypothalamus of the brain, is specific to filtering out visual noise and to accentuating edges. Adjacent the dLGN is the Medial Geniculate Nucleus (MGN), which is specialized for tuning and accentuating auditory signals (5).

A typical property of the brain is that of parallel and hierarchical processing (2). Multiple structures of the brain may operate at the same time, yet are also linked within a hierarchy. In the visual system, a visual image is supplanted from the 3D world onto our 2D retinae. The retina is covered in photoreceptors which are tuned for specific wavelengths and intensities inherent in the visual stimulus (2). These receptors, working in parallel, send their signals to the next stage of processing within the retina, where both spatial and temporal properties are encoded. These signals fall upon "ON-center" and "OFF-center" ganglion cells, each of which responds in a particular way to the pattern of light falling upon their receptive fields (2).

Note that there are many additional retinal layers located between the ganglion cells and photoreceptors, and all of these cells do not produce action potentials. The task of many of these cells is to further help localize edge components spatially, as well as to accentuate these components. Additionally, another set of these retinal cells serve to signal changes in time. That is, when an object is moving in time, these cells will help detect the changes occurring within the visual field (2) (5).

The signals generated by the ganglion cells are further sent to the bilateral dLGNs; this wiring is very important. The visual field covered by the retina is such that the temporal fields route to the ipsilateral dLGN, and the nasal fields route to the contralateral dLGN. Because primates have eyes which both face forward and have some overlap within their visual fields, there are also binocular cells within the dLGN. A binocular cell is a cell that responds to a stimulus shown in either the left or right eye, or responds to the area where the two visual fields from the two eyes overlap (2).

The dLGN contains 6 layers comprising at least two systems of cells, magnocellular and parvocellular, each tuned for different properties of the visual scene (2). The magnocellular system responds best to properties such slight differences in contrast, low spatial frequencies, and is sensitive to luminance; it comprises the two deepest layers of the dLGN. The parvocellular system responds best to details, high spatial frequencies, and color; it comprises the other four layers. These layers alternate with respect to which eye the incoming signals originate The main job of the dLGN is to accentuate the contrasts of the visual stimulus and to reduce noise.

At the level of the dLGN, signals from the two hemi-fields remain separate, converging later at the level of the cortex. Through the optic radiation, signals enter through Layer IV of cortex, which is the entry point for all signals into the primary sensory areas of the brain (5) (2).

In the visual system, this area is called “Primary Visual Cortex” (also known as V1 or “striate cortex”). This is the first layer of signal processing. Primary visual cortex is where the orientation components of a visual scene are parsed and encoded (5). For example, a picture of a square will stimulate both horizontal edge-tuned cells and vertical edge-tuned cells. This is because a horizontal edge-tuned cell will fire when the code for a “horizontal edge” is present and is silent otherwise. These orientation signals are then sent further upstream where, together with input from the cognitive centers of the brain, they are put together into perceptual hypotheses regarding what is being sensed (and what is ultimately perceived).

2.1 Features of Primary Visual Cortex

Over 50 years ago, David Hubel and Torsten Wiesel explored the properties of a mammal’s visual circuitry by carefully inserting micro-electrodes into the primary visual cortex (V1) of anesthetized cats. Through their work, they introduced the first model of V1 known as the Ice Cube Model (9). The model proposes that for a given position in space, a cortical column in V1 is tuned for properties of orientation and ocular dominance. While this model does account for the cellular arrangement of V1, it does not accommodate for other properties in which these cortical cells are tuned, such as spatial frequency or color processing.

Years later, activity in V1 could be marked using 2-deoxyglucose, a substance readily taken up by active neurons (10). This visualization technique confirmed Hubel and Wiesel's Ice-Cube model of V1, yet it also revealed cells arranged around the blobs in pinwheel-like swirls; gradients of orientation preference that were "clockwise" or "counter-clockwise" (11) (12) (13) and aligned along the edges of the orientation columns.

In particular, 2-deoxyglucose staining revealed two key features of V1 organization. Not only were there pinwheels (singularities), but also large swaths of cells all with the same preference for an orientation (iso-orientation domains). The iso-orientation domains came together at the singularities, and in different species these iso-orientation domains can have a characteristic shape and organization.

2.2 Current Models of Primary Visual Cortex

Many models attempted to explain the self-organizing behavior of V1 cells given visual experience, as well as to explain the rules by which they organize into singularities and iso-orientation domains. The three layer approach is considered to be biologically sound, meant to represent the hierarchical structure of the retina, dLGN, and V1 cortex progression of cellular organization (14) (15). Bhaumik and Mathur used this model within the context of a purely feed-forward approach, and using Hubel and Wiesel's idea that complex cells were a simple, linear summation of simple cells. A simple cell is a cortical cell whose receptive field can be mapped as there are specific excitatory and inhibitory areas. These cells respond strongest to light stimuli that are the same shape and orientation that fall directly in line with their field mappings. On the other hand, a complex cell does not have a mappable receptive field and instead responds to any stimulus (light or dark) that is the preferred shape and orientation for the cell, and is moving in a particular direction (2) (9) (16).

Such an approach is plausible because it was not necessary to pre-program the shape of the cell's receptive fields into a difference of gaussians formation; they found that the formation of ON-center and OFF-center areas of the retinal ganglion cells happen "naturally" and that the resulting projections from these areas contributed to the formation of orientation preferences downstream in V1. Further, when incorporating non-linear contributions from neighboring cells (17), cells in V1 show an even finer tuning of orientation selectivity, which is also true for what happens within V1. This natural progression of such properties is a characteristic feature of brain development (5).

If a cell is binocular, then it can be stimulated by a signal from either of the two eyes. For example, a stimulus shown to the right eye or the left eye causes the same cell to fire. As a result of this property, the cell is classified as binocular. This binocularity plays an important role in the organizing of V1 topographic maps. In 2-deoxyglucose stained maps, pinwheels are typically found at the edges of binocular columns whereas iso-orientation domains are found more centrally (18). Tuning influences from neighboring cells, as well as binocular contributions, are thought to play a role in the development of orientation selective cells in V1.

Erwin (19) provides an excellent review of the most salient models of orientation selectivity in V1 and notes the general properties and assumptions of most models of V1.

First, many assume neurons to be organized in 2-dimensional sheets, where the receptive fields of these cells are represented by either feature vectors or as synaptic weight vectors. Also, non-linear contributions by complex cells, as well as feedback from intracortical or higher areas, are also represented in some fashion (unfortunately, these are typically programmed directly into the model). Almost always, most models incorporate two biologically intuitive properties: continuity and diversity. Continuity is where nearby columns of cells in the cortex tend to prefer stimuli with similar features (via the average or convolution of a defined area across a given period of time), and diversity is where the space of all feature preferences ought to be filled in as completely as possible (demanding that all of the "neurons" eventually form an orientation preference). This is typically done with band-pass filters or competitive networks, yet it is not fully known whether the brain operates in either of these to ways. Additionally there is the property of "global disorder", meaning there ought to be some noise added into the signals creating the map, as noise is inherent in any biological system.

2.3 Current Problems in V1 Modeling

Previous developmental models have been able to reproduce the overall organization of specific feature maps in V1, such as orientation maps, but are generally formulated at an abstract level that does not allow testing with real images or analysis of detailed neural properties relevant for visual function.

For example, it is commonly accepted that a Self-Organizing Map is a biologically relevant way to create clusters. Self-organizing maps (SOMs) use a learning technique called Konohen's algorithm, which relies on unsupervised learning to organize similarly preferring neurons into local sub-networks. Swindale and colleagues have found success when using SOMs in the organization of cells in V1 (20) (21). However, despite his success in generating seemingly good biologically similar maps of V1, his single layer approach leaves room for improvement. Additionally, whether or not the biological brain utilizes a SOM algorithm is not known, and probably unlikely given that such cellular behavior patterns have not been recorded empirically. Such models make assumptions that may not agree with biological function. Whether or not the brain actually implements a self-organizing mapping algorithm has not been confirmed.

The model within this thesis attempts to make no such operational assumptions about how cells fire when within a network; it simply assumes the properties of a neuron unit and allows for any network-like properties to develop naturally when these units are placed in a network.

CHAPTER 3

MODEL ARCHITECTURE

It is important to build an architecture that agrees with biology. Thus, the model presented here is built with a biologically plausible architecture in mind using biologically plausible units. The model follows two main tenants known in the organization of the brain: parallel and hierarchical processing.

3.1 Parallel Processes: Leaky integrate-and-fire neuron

This model is built upon the neuron model proposed by McCulloch and Pitts (22) and is shown below.

$$y = f(\sum_{i=1}^{n} w_i x_i)$$

Where f represents the step function or threshold for firing (it can return a 0 or 1), n is the number of dendrites belonging to the neuron, x is the input signal for dendrite i, and w is the weight given to the signal on the dendrite (see Figure 2).

The function f returns 1 if the membrane potential is greater than -35mV and returns 0 if anything less. The resting potential is then reset to -70mV. These are the values found in biological data (5).

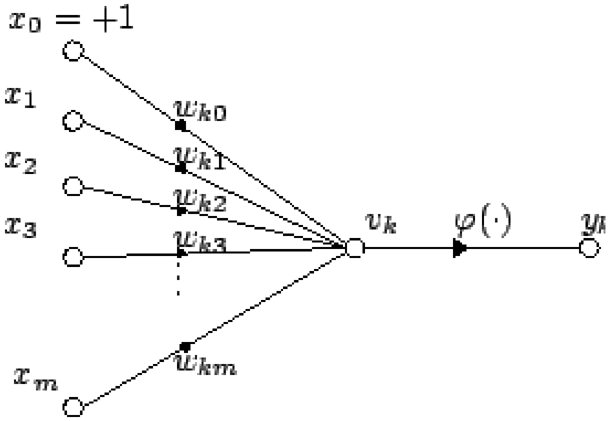

Figure 2. Artificial Neuron Model
x represent incoming dendrites, which carry signals. w are the weights on each dentdrite for neuron k. v is the sigmoid function of the weights and dendrites. $\varphi(.)$ is the step function. y is the output.

3.2 Hierarchical Processes

There are three main hierarchical processes present in the model. These processes are from retinal "photoreceptors" to the retinal ganglion cells, the retina to the dLGN, and from dLGN to V1.

3.2.1 Photoreceptors to Ganglion Cells

Images are shown to a retina containing "photoreceptors". The human retina contains 126 million photoreceptors, 120 million of those being the black-and-white detecting (non-color selective) "rod" photoreceptors, the other 6 million being color-selective "cone" photoreceptors. These receptors are part of the receptive field of ganglion cells, which are the signaling units to areas outside the retina. There are approximately 1.5 million ganglion cells in a single human retina (5).

The human fovea covers approximately 2-degrees of the visual field, comprising approximately 108,000 photoreceptors. In the primate fovea (presumably including human) the ratio of ganglion cells to photoreceptors is about 2.5; almost every ganglion cell receives data from a single cone, and each cone feeds onto between 1 and 3 ganglion cells (5).

Because I am only interested in the development of orientation selectivity and not acuity, this model will not consider inputs from the fovea. However, the fovea is used as an anchor when selecting which retinal efferents will map onto the next stage of visual processing. In neuroscience, visual hemi-fields are divided with respect to fixation which in experimental research, is always placed within the fovea.

The photoreceptor layer simply relays a binary signal to its ganglion cell(s). A photoreceptor fires a 1 if the (x,y) location is less than or equal to 128 RGB, and does not fire if the (x,y) location is greater than 128 RGB. All images in the dataset are black and white without shades of gray.

The ganglion cell (a neuron unit) receives the incoming signal and fires if it receives enough signals to overcome its threshold. These cells fire most robustly when it's neighbors are not firing, a property known as "lateral inhibition". Lateral inhibition allows for contrasts (and edges) to be detected more robustly because with this set up, ganglion cells will fire where large contrasts within the visual stimulus exist, the first stage in signaling the presence of edges (2).

3.2.2 Retina Layer to dorsal Lateral Geniculate Nucleus Layer

Signals exiting the retina are carried by the ganglion cells. However, before the signal can reach the dLGN, it travels through the optic chiasm and bifurcates as seen in Figure 3.

The ipsilateral fibers receive signals from cells responsible for covering central vision (which are mostly binocular) and the contralateral fibers are responsible for those cells covering peripheral vision (which contain both binocular and monocular cells). It is important to maintain this distinction because the organization of these fibers influences the ultimate organization of cortical neuron orientation preferences.

Once the signals reach the dLGN, edge components are further accentuated by applying Hebbian learning. It is known that cells that fire together will wire together (5), therefore at the level of the dLGN, cells that fire together on the same time steps are likely signaling an edge. For example, if a cell's neighbor immediately above and immediately below fire in the same time steps, the three of them will reinforce each other's firing, and are now more likely to fire together if any single one of them receives input. This set up is likely the beginnings of a vertical edge. The cortex is left to interpret these signals, but at the level of the dLGN, Hebbian wiring is already taking place.

3.2.3 dLGN and Binocular Cells

Binocular overlap is the result of two overlapping receptive fields. Humans and other animals that have a pair of eyes facing forward have binocular vision. Binocular vision allows for depth perception, an important part of the visual experience. Other animals have eyes on the side of

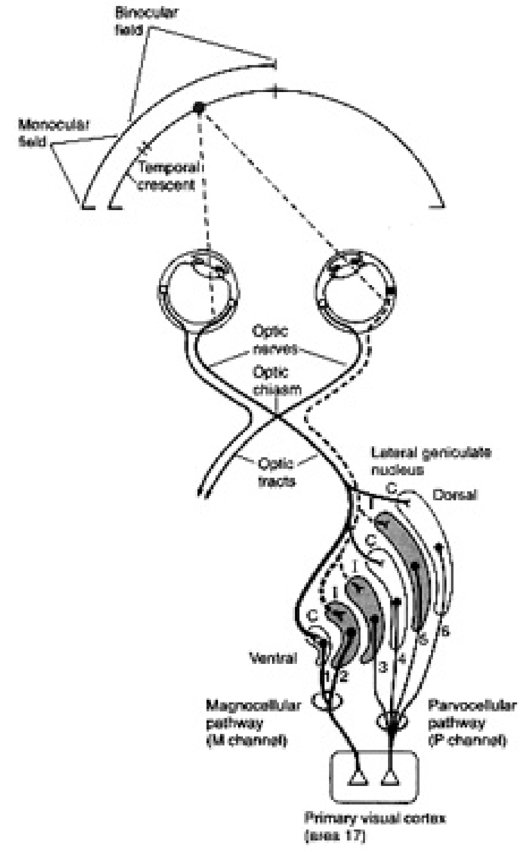
Binocular field
Monocular field
Temporal crescent
Optic nerves
Optic chiasm
Optic tracts
Lateral geniculate nucleus
C
Dorsal
I
C
I
C
6
5
4
3
2
1
Ventral
Magnocellular pathway (M channel)
Parvocellular pathway (P channel)
Primary visual cortex (area 17)

Figure 3. The Visual Signal Cascade
A diagram showing how input from the visual field is received by the eyes and travels through the visual system of the brain. This diagram shows the signal moving from the receptors in the retina, where each hemi-field has fibers connected to a particular dLGN. At the level of the dLGN are 6 layers, where two cellular systems reside. These signals exit the dLGN and enter the Primary Visual Cortex.

their heads. These animals do not have depth perception, but in exchange they have a wider field of view. Binocular cells receive a signal from each eye meaning that both eyes can see it (see Figure 4).

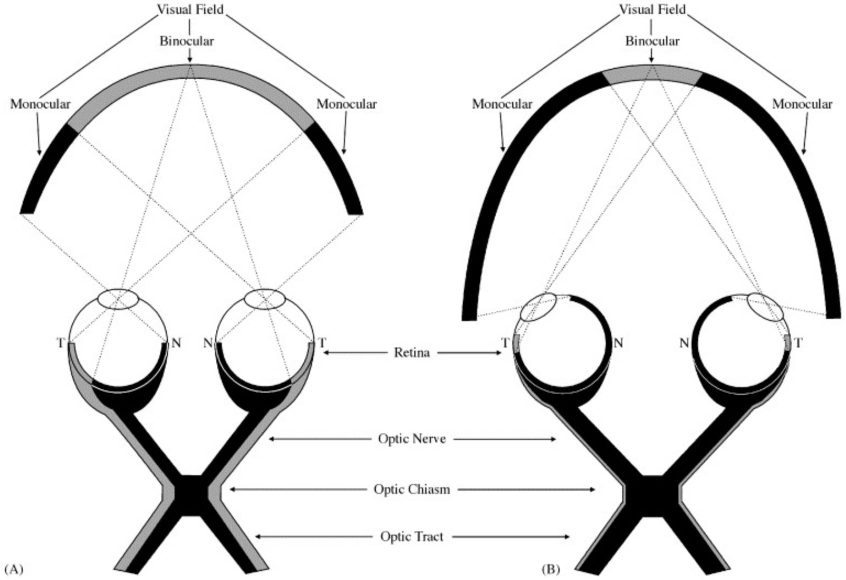

Figure 4. Binocular overlap of a pigeon and a human
Left: visual field of a human. Note the large amount of binocular overlap between the eyes.
Right: visual field of a pigeon. Note the larger visual field.

For example, in Figure 4, we can say that the visual field is sized 200 degrees (monocular and binocular portions summed). The binocular overlap is 120 degrees (the monocular portion). The image at time step t is projected onto the entire space (the monocular and binocular portions). Therefore, in the model, the left eye receives the portion of the image from the leftmost monocular portion plus the binocular portion, and the right eye receives the portion within the binocular portion plus the rightmost monocular portion. This results in the two eyes receiving nearly the same image, offset only by the amount of monocular vision. The more binocular overlap, the more similar the images.

In the model, signals from both the left and right eyes contribute to a binocular cell's firing threshold. If the signals are different, then the cell looks at its neighbors to check which orientation is preferred in the near vicinity. The monocular cells receive signals only from their respective retinal ganglion cells. The dLGN in this model contains only one layer for orientation selectivity. In the biological brain, these signals remain separate until they leave the dLGN, but because I am not considering the other dimensions of the visual scene, adding them in dLGN as I do fits well with biological processes.

3.2.4 Primary Visual Cortex (V1)

V1 will receive inputs from the ipsilateral dLGN. In the human brain, V1 is separated by each visual hemisphere (the signals do not converge until later in the hierarchy). Therefore, as the network develops, it keeps each each hemisphere of the visual scene separate. By keeping the two separated, we can then do further experiments by blocking the signal cascade at one particular point and observing how it affects the output. After the signals from each dLGN come together, the entire map is created by combining the two maps (one from each eye) by placing them next to each other using fixation as a reference. The result should be somewhat symmetrical about this fixation point.

Additionally, the cortex comprises six layers with a few of them participating in recurrent feedback. The cortical layer of this model has neurons that accept inputs from many dLGN cells. Additionally, the dLGN cells send their signals to more than one ganglion cell. Both of these features are biologically accurate. This part of the model allows for the orientation preferences to be "smoothed out" so to speak. This overlap of both firing and non-firing cells allows for "ON" and "OFF" areas to form in the cortex, resulting in the equivalent of a "blurring" of the orientation preferences. The cells on the edges of the iso-orientation domains receive influence from many neighbors as well as the afferents of their neighbors, thus the ultimate orientation preference of these cells is determined by both the neighbor's preferences (parallel) and the arriving signals from the dLGN cells (hierarchy).

Lastly, lateral inhibition in the cortical cells helps to "sharpen" the blur due to the interconnections between the dLGN and the cortical cells. If a cell receives signals but its neighbor does not, then lateral inhibition will allow an edge de-marking the boundary between two iso-orientation domains to develop, if the connections are right for it to happen. The development of these boundaries is entirely determined by the influence of the neighbors and the strength of all signals (both lateral and afferent) coming into the cell.

CHAPTER 4

MODEL PARAMETERS

The basic architecture of this model is inspired by the biological brain. It consists of three main layers, meant to parallel the retina, LGN, and cortical layers of the biological brain. Each layer is composed of a sheet of leaky, integrate-and-fire neurons. Importantly, the input/output relationships between layers is typic to the biological brain; that is the layers are wired as they would be in nature to the best of the current knowledge.

Each of the sections below summarizes a parameter of the model. Note that none of these parameters change the functioning of the network directly, rather they change how a single neuron unit responds. Any network behavior is simply the emergent result of the parameters followed by the interconnected neuron units.

4.1 Image Dataset

Images presented on the model's retinae are pulled from an Image Dataset. This image dataset contains 100 images, as well as duplicates of each of these images, rotated every 5 degrees, ranging between 5 and 180 degrees. The 100 images are the numbers 0 through 9, the English alphabet in both uppercase and lowercase letters in font Ariel, as well as various geometric shapes such as circles, squares, or interesting textures and high contrast pictures from across the internet containing significant numbers of edge components. In total, the Image Dataset contains 3800 images. See Figure 5 for an example:

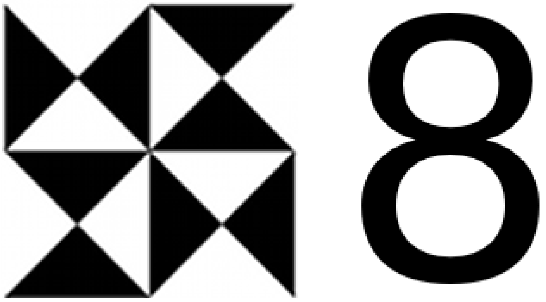

Figure 5. Examples of training images
On the left is an example of a geometric shape with edges, on the right is the number 8.

4.2 Time steps

In this model, a time step is synonymous with a single presentation, and they are always equal. Therefore, on each time step, a single image is shown to the model's visual field. For example, as the human visual brain develops, a variety of visual images are nearly continuously being projected onto the retina. In this model, a single time step is akin to 200ms, or when comparing to the human visual system, one glance (known as a "saccade").

The number of time steps in the model can be modified. This changes the number of images that the model is exposed to and thus, which images influence the "cortical" development. Fewer images will be less effective than more images. Because this model follows that cells of the same orientation preference will likely be found closer to each other, later time steps will serve only to tune the orientation preference of cells, not modify them. Therefore, a rough model of visual cortex can be generated by only a few images, but for a more accurate model, more time steps will be required. There may also be a ceiling to number the of time steps such that further images will have almost no influence on the orientation preference of the "cortical cells", especially given the `reinforcement` variable (discussed below).

Additionally, the size of the subset of images can be controlled. For example, I could begin an instance of `t`=100 time steps, but I would want to limit my sampling pool of images to 10 images. The images shown to the model are always randomly presented, but one is able to control the size of this sample. This allows for investigation of whether or not a varied set of input images influences the development of orientation specific cells. Based on the experiments of Hubel and Wiesel (23), one would hypothesize that by limiting exposure to an enriched set of images, one would see diminished development of certain orientations. Specifically, if the set did not contain many images containing edges at 45-degrees, it is likely that not many cells would be selective for 45-degree edges.

4.3 The Retina

In the model, the dimensions of the photoreceptor sheet can be modified. However, in the human fovea there are approximately 100,000 photoreceptors. This makes the dimensions of the retina approximately 330x330. The retina size run in the simulations in the demonstration of this model is a fraction of this at 100x100.

4.4 Binocular Overlap

This model allows for the degree of binocular overlap to be modified. If I wanted to create maps of human visual cortex, I could create a visual field that was 200 degrees total, 120 degrees of which overlapped. On the other hand, if I wanted to create maps representing pigeon cortex, the visual field size would be 400 degrees of which only 40 degrees overlapped. By modifying this parameter, I can test to see the extent to which binocular overlap affects the resultant orientation preference map in the cortex. See Figure 6 for an example of the input images given binocular overlap.

4.5 Starting Weights

Primates and cats are born with pre-defined orientation selectivity even before the animal opens its eyes. These orientations are later modified by experience (24) (25). However, if these pre-defined orientations are weighted heavily, it would take more time steps in order to change the orientation selectivity of the cell in question. Therefore, in the model, one can modify the starting weights of the orientation preference of the cells. It is presumed that with a heavier starting weight, more time steps are required to "overcome" the influence of the starting weight on orientation preference.

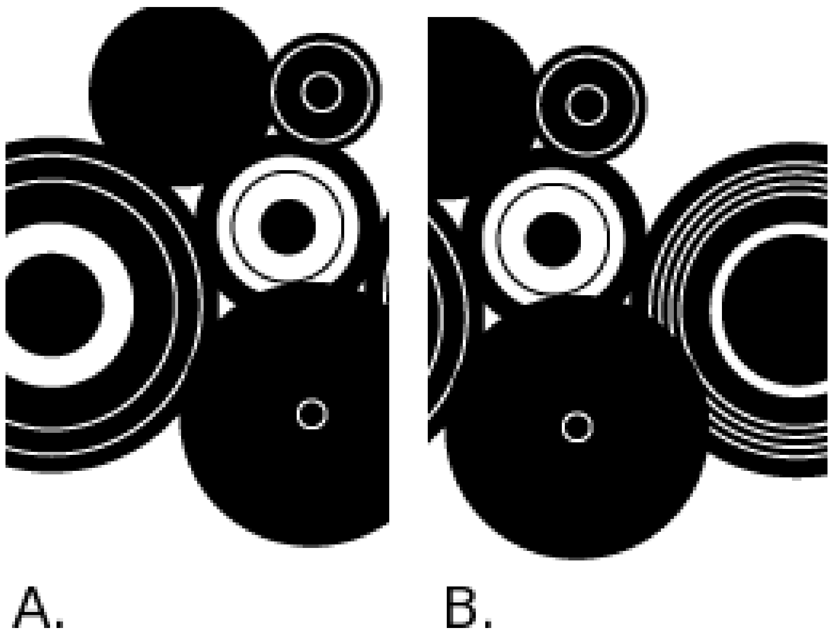

Figure 6. Training Images Example
A. Image shown to the left eye. B. Image shown to the right eye. Note some areas overlap (binocular areas) and other areas do not (monocular areas)

4.6 Iso-Orientation Domain Size at t=0

One fact that is implemented in building models of V1 is that of an already established "starting map" that can be mapped before the animal experiences the visual world. Following the rule that cells with a similar orientation will likely reside together spatially, one set of these cells is called an "iso-orientation domain". How large these swaths of cells are at the beginning could influence the resulting size of the iso-orientation domains. It is assumed that only the cells on the "fringe" might be influenced by the neighbors on any particular iteration. Then, as the time steps progress, the iso-orientation domain area might shrink, until the threshold level is unaffected by the influence of the neighbors. Hence, it is assumed that the larger the iso-orientation domains are at the start, the larger they will be at the end.

Iso-orientation domains are initialized in the model by randomly placing squares of a parametrized size about the "cortex sheet", each with a randomly selected orientation preference assignment. The number of these squares that are placed is the same as the retina width, which guarantees that these areas will overlap. Any other area in cortex that has not yet been covered by the random placement of the domains is filled with a random orientation preference. These areas should evolve with respect to their neighbors.

4.7 Threshold

Many cells of the nervous system operate in terms of thresholds. A cell either fires or it does not fire: the code is in the frequency. Therefore, all cells in this model also operate by thresholds. In the cortex, a cell has a threshold whereby it will change its orientation preference.

Meeting this threshold depends upon the orientation preference of the neighbors. Essentially, the current cell pings the other cells falling within its immediate vicinity (8 neighbors) and checks their orientation preference. If the number of cells within the vicinity prefer a particular orientation that is different from the current cell, then the current cell will change its orientation to match. This threshold value can range from 1 to 8. It is assumed that a threshold that is too low will now allow for domains to form, and a threshold that is too high will force the map take on a single orientation preference.

4.8 Reinforcement

Additionally, when a cell receives a signal about a particular orientation, there is a set amount of "reinforcement" for learning this orientation. This reinforcement value competes with the threshold value. The difference between the current cell and its neighbors does not play a role in the threshold of the current cell before it makes the switch; it is winner-take-all for the "winning" orientation. Independent of the current cell's orientation preference, if the neighbors override the current cell's orientation preference, the current cell will adapt and take on the majority orientation preference.

If the reinforcement value for an orientation is large, then when it comes to modifying the current cell's orientation preference based on the neighbors, the contributions from the neighbors must be very large to "offset" the reinforcement value the current cell received for learning that orientation. Such a case would occur if for x time steps in a row, a cell received the signal for a horizontal preference. If after x time steps, the other cells are of vertical preference, the reinforcement for the current cell to remain horizontal might override the influence of the neighbors. This would especially happen if reinforcement values are too large.

4.9 Arborization Factor

Lastly, within the model's cortex, the arborization factor can be changed. This is the degree of dendritic spread that the dLGN cells use to connect with the cortical cells, as well as the amount of inputs the cells of the cortex receive. This value is based off of a Gaussian (Gabor-like receptive fields), and the standard deviation of the Gabor can be modified. The receptive field shape used is known as the Difference of Gaussian's model, a very common model used to describe receptive fields of cells within the visual system(2).

4.10 Output

The final output is a color-coded map containing 6 colors. Each color represents an orientation. While visual cortex can detect orientations along a continuous scale of 0-180 degrees, I have divided this range into 8 discrete "bins" : 0, 30, 60, 90, 120, 150. The resulting map is color-coded by orientation preference according to these 6 colors.

I expect these colors to form iso-orientation domains which are connected via interspersed "pinwheels". However, an area is not considered a "pinwheel" unless all colors meet at a point, representing all orientations between 0 and 180 degrees.

A pinwheel can be clockwise or counter-clockwise. This nomenclature refers to the direction that the orientation preference of the cells goes when moving about the center of the pinwheel.

For example, a pinwheel where the order of the orientation preference moves from 0 to 180 occurs clockwise, then this pinwheel is clockwise. On the other hand, if it moves from 180 to 0, then it is counter-clockwise.

4.11 Step by Step Model Summary

Let us visualize this model as a basic system. The input for this system comprise a series of black and white images into the two “retinae”. On each time step, a different image is projected onto the retinae. The “receptors” of the retina “fire” if light stimulates them or are silent if they are left in the darkness (“1” if a white portion of the image; “0” if a black portion of the image). These signals are sent to the ganglion cells, where differences between the photoreceptors are compared. If a ganglion cell contains two photoreceptors, one signaling “1” and the other signaling “0”, its electrical potential will become more depolarized. If there are enough cells with different signals within the ganglion cell’s receptive field, then the electrical potential of the ganglion cell will be met, and the cell will fire off an action potential.

The signals from the ganglion cells are sent to the dLGN, but not before they pass through a very important wiring. Cells whose receptive fields are in the nasal hemi-fields route through the optic chiasm, effectively crossing to the other side of the brain. Cells whose receptive fields are in the temporal hemi-fields remain on the same side of the brain. Both sets of fibers meet at an dLGN, but the signals remain separate.

Once inside the dLGN, the receiving signals begin to detect edge components. The ganglion cells that detect differences within their receptive fields are signaling change, and thus, it is in this location where an edge is likely present. For example, if the receptive field is a 3x3 grid, and all cells in the third row fire on the time step, then the ganglion cell will signal a change has occurred in its receptive field. dLGN cells also take note of their neighboring cells and will also actively inhibit all neighboring cells around it if it fires. This allows for edges to be accentuated.

These signals come together as they leave the dLGN and feed onto the primary visual cortex. At this point, cortical cells look for dLGN cells that are firing in particular orientations. If of 9 cells in the receptive field, all cells in the second row are firing, then the cortical cell will begin to prefer horizontal edges. It will compare with its neighbors, which have overlapping receptive fields, to see if they have also detected similar components in the area. If so, they reinforce each other's firing. If not, the cortical cell may change its orientation preference to match its neighbors. It will not change its orientation preference if either not enough neighbors are of the same preference, or if it has been highly reinforced for a particular orientation (which could happen if the particular cell was shown only horizontal edge components in the entire time sequence).

The cells will update their thresholds and preferences with each iteration, where each iteration is a new time step. At the end of all time steps, a map of orientation preferences is generated by color-coding the preferences of the cells at the last time step.

CHAPTER 5

MODEL EVALUATION

In order to evaluate the model, first we must understand what we can and cannot control, and then we must be able to compute some quantitative measures on a set of maps and compare these quantitative measures across both model maps and empirical (real) maps. In order to do this, we must make sure that these measures are not dependent on receptive field size as empirical maps are not created using a standard scale. For example, a receptive field in one lab's empirical map might be twice as large as another lab's receptive field.

5.1 System Summary

Figure 7 diagrams the V1 model system. Gray areas are hardwired whereas white areas can be controlled.

5.1.1 Input

The input are the images projected onto the retinae. One can control the number time steps, number of images, and the pool from which the images are chosen. For example, for the normal "brain", one would provide an enriched environment with a variety of images. If interested in development without horizontal edges, one would restrict the images shown to the model by providing it images without horizontal edge components.

5.1.2 Retina Parameters

The retinal ganglion cells receive signals from the receptors. One may change the *size of the retina*. There is no fovea because the model cannot move its “eyes” around.

Additionally, *binocular overlap* is a parameter that can be modified. If interested in how animals with good depth perception develop, binocular overlap would be set high. If interested in how animals with no depth perception, but a large field of few develop, binocular overlap would be low.

5.1.3 Cortical Cell Parameters

The *starting weights* of the cells can be set. They will be modified by experience, but their starting values could have an effect on the outcome, depending on how many iterations.

The *size of the iso-orientation domains* can also be set. The larger the orientation domains then the larger they should be at the end.

The threshold for changing orientation preference. One can decide on how the weight of the influence of the cell’s neighbors on deciding orientation preference.

Reinforcement is the additional weight added to a cell’s particular orientation preference.

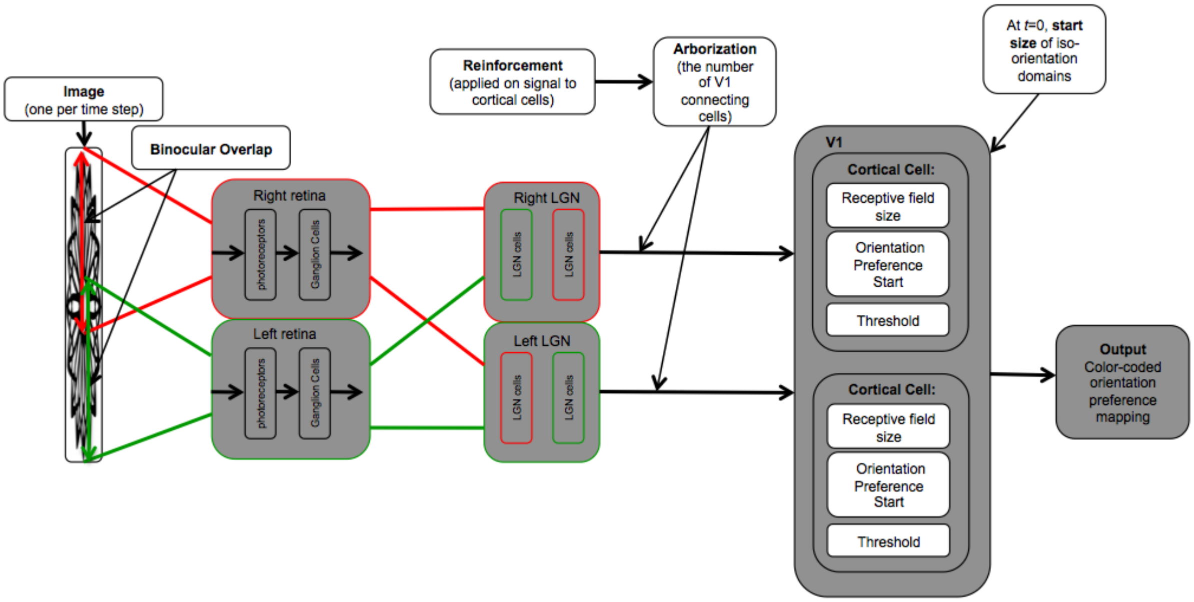

Figure 7. The V1 System
A diagram of the system, indicating the changeable input parameters (white) and the portions that are not modifiable (gray).

5.2 Minimum Spanning Tree

At this point, we have a color-coded orientation preference mapping. We can place this map on a cartesian plan, and given the locations of each pinwheel or iso-orientation domain centroid, we simply treat each as a vertex in a graph. For each set (pinwheels or iso-orientation domain centroids) a minimum spanning tree (MST) is generated using Prim's algorithm (see Figure 8). Edge weights were the Euclidean distances between any two vertices.

Prims algorithm is an algorithm that generates the MST for a connected, weighted, undirected graph. The tree that it forms includes every vertex, where the total weight of all the edges in the tree is minimized, therefore Prims algorithm guarantees that all vertices are contained within a single, fully connected graph; each map has a unique MST. We constrained the implementation of our algorithm to allowing each vertex a maximum valency of three. This is important because it allows each map to have a unique fingerprint based on the input parameters (and the model itself) that generated it.

5.3 Coefficient of Variation

The resulting collection of edges within each set (pinwheels or centroids) for any given graph allows us to perform our measures. The first measure is the coefficient of variation (CV), a dimensionless coefficient defined by the standard deviation of the edge weights (length) over the mean edge weight for a particular graph. Using a dimensionless coefficient allows us to ignore receptive field size or the resolution under which a topographical map was generated. We can also examine the CV of the MST edges of only the clockwise pinwheels and compare

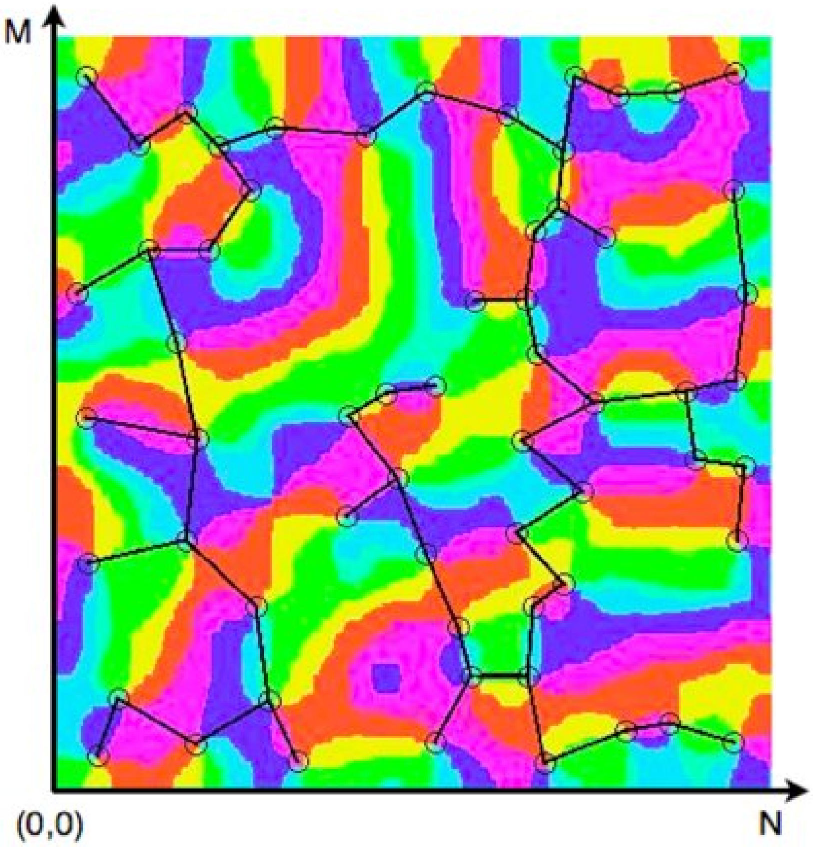

Figure 8. Example of MST on map
A minimum spanning tree is computed on a map using pinwheels as verticies.

this to the CV of only the counter-clockwise pinwheels to get an idea of the distribution of pinwheels depending on type. The CV of the centroids of each of the iso-orientation domains is also computed.

5.4 Ratio of Counter-clockwise to Clockwise pinwheels

The ratio of counter-clockwise to clockwise pinwheels is also a measure of key interest. A clockwise pinwheel is one where when moving clockwise about the center, the progression of colors correspond to the increasing gradient of orientation specificity, and vice a versa for the counter-clockwise pinwheels. Examining this ratio is a measurement already in practice, and thus it is included in these analyses.

5.5 Compactness of Domains

For the iso-orientation domains, not only did we statistically compare the edge weights between centroids, but we also examined the compactness of each iso-orientation domain, or the "stringiness" of each (or how not compact they were). Compactness is a dimensionless measure dened as the perimeter squared over the area. To differentiate noise from true iso-orientation domains, we altered these sections such that any area comprising less than 0.05% of the total area of the map was discarded.

5.6 Model Validation

In order to validate the maps generated by my model, all evaluation measures are compared to those measured in the real data. A dataset containing the CV of the pinwheels and domains, and the CCW/CW ratio is analyzed by its principal components. In general, this measures the degree to which these parameters are correlated, and any datasets that overlap are considered highly similar. Therefore, if the measures obtained by the generated and empirical maps appear in the same region of PCA space, then they can be considered biologically similar.Before we can validate the model, first we must preprocess the biological maps in order to make appropriate comparisons. First, each map is standardized through remapping the original color encoding (256-color map) to eight standard colors. First, we apply global histogram equalization to each image so that, for example, all shades of red are mapped to the same value, and so on. A distinct color-code (colors red, orange, yellow, green, light green, blue, violet, pink) is used to represent each type of the orientation-preferring cells (0, 22.5, 45, 67.5, 90, 112.5, 135, and 157.5 - degrees respectively); if only 6 of these colors represent cells preferring orientations 0, 30, 60, 90, 120, 150). Each map is already discretized and so the remapping is performed to rid each image of noise as the result of image compression.

Given each normalized, 8-bit color encoded image map, we set each image on a 2-D Euclidean plane ascribing the lower-left corner the coordinates (0,0). Scanning the image, we collect the (x, y)-location of each pinwheel and iso-orientation domain centroid. A pinwheel is defined as an area within a standard radius where all eight colors had to appear in a clockwise or counter-clockwise order. The iso-orientation domain centroid was defined as the center of gravity of the shape created by the domain.

After preprocessing, all empirical maps are processed and the same measures as described above are obtained.

CHAPTER 6

RESULTS

The output of my model is a color-coded map of orientation preferences within primary visual cortex. I have evaluated each of the maps according to the evaluation measures described in Chapter 5. Now, I must make comparisons to test the power of my model. In this section, I demonstrate the parameters settings which produce statistically similar maps, and also show which maps should not be considered biologically similar.

6.1 Principal Component Analysis

Principal Component Analysis is a way to find out which components of a dataset can account for most of the shared variance within the set. Using the first two main components, it can demonstrate the extent to which multiple variables are correlated, and the extent of their relationships. In this thesis, a PCA can tell us the relationships between the input parameters as well as tell us the relationships that the evaluative measures have in the both model maps and empirical maps. I demonstrate a PCA using both the first two and first three components to locate cluster formations.

For example, to find the first principal component, I would draw a regression line through all of the data in multidimensional space (n dimensions, where n is the number of factors or variables). This is my principal component. It is the line of best fit, and will serve as my X-axis.

My second component is a line drawn orthogonal to the principal component. It will capture any variance in the orthogonal direction to the principal component. Then, these two lines are rotated such that they can be visualized as an X and Y axis. These two components capture a good majority of the variance, yet can be extended to adinfinitum.

6.1.1 Input Parameters

Figure 9 shows the PCA of the input parameters. The figure shows that the variables TimeSteps, Samples, and RunTime are all highly correlated for this set. These variables are correlated because the more time steps required, the more likely a higher number of images will be sampled, and this also takes a higher amount of time. Orthogonal to these variables is StartingWeight, and then opposite of that are variables Threshold and Reinforcement. StartingWeight is negatively correlated with Threshold and Reinforcement, meaning that the higher the starting weight, then the lower Threshold and Reinforcement are in this dataset. Lastly, the variable StartingSize is orthogonally correlated with the other variables. All other variables were held constant for the data points presented in this thesis.

Biologically, these parameters are meaningful. Timesteps, Samples, and RunTime would also be well correlated; the more images the animal is exposed then the more time it takes to view all of those images. If I show an animal 10 images, then that is ten time steps, which will take a shorter time to run than showing an animal 100 images, requiring 100 time steps.

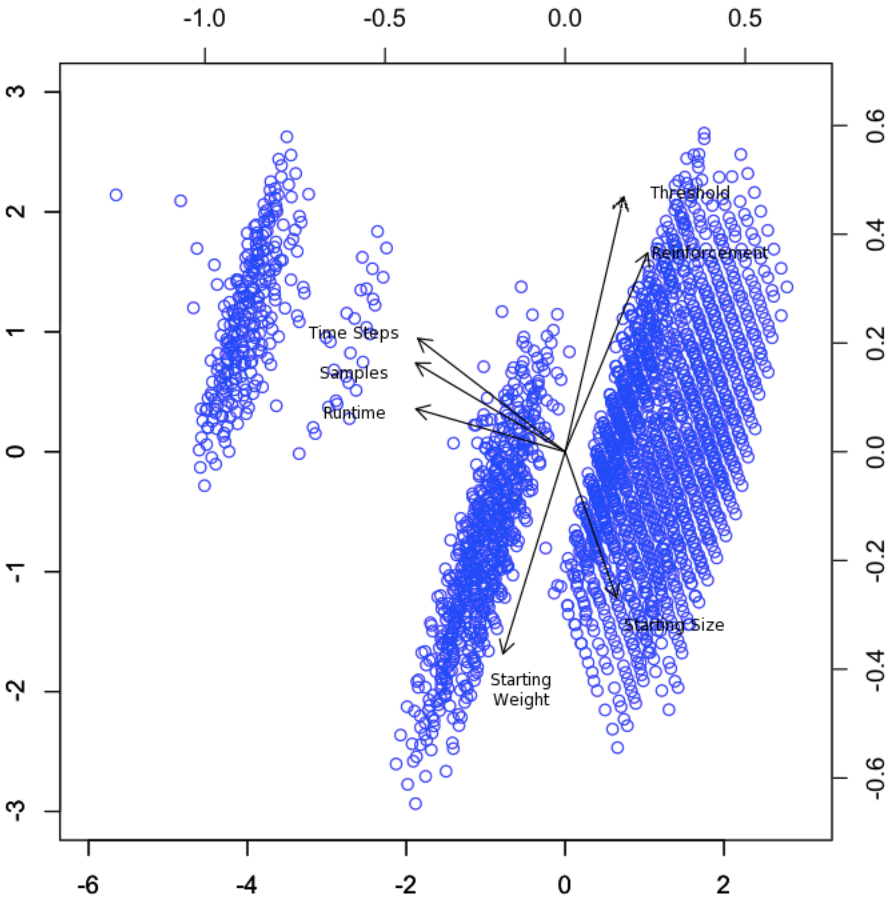

Figure 9. PCA of main input parameters
This PCA analysis shows four distinct clusters for this set of data.

Startingweight is biologically similar to the findings by Hubel and Weisel where they found kittens, who had not yet opened their eyes and thus had not started to view any images, already had V1 cells with orientation preferences, albeit they were very weak.

Negatively correlated to startingweight are threshold and reinforcement. These are two properties of the cell itself. Threshold is the amount of neighbor influence. In biology, a cell actively inhibits its neighbors. If a particular cell is being inhibited by many neighbors in a single time step, then it will likely change its firing to sync in with the neighborhood; that is if a cell receives inhibitory signals from all of its neighbors, the inhibition from the neighbors is much stronger than the inhibition from the single cell. As a result, the single cell will no longer be reinforced for its original orientation, and will begin to fire with its neighbors. Typically, cells that fire for similar stimuli are also spatially nearby.

Also, depending on the frequency of the stream of action potentials arriving at a cell, the reinforcement for a particular wiring can be weak or strong. If reinforcement is high, then high frequency pulses "wired" the two cells together (Hebbian Learning). In the PCA, these two variables are positively correlated because, to quote Donald Hebb, "cells that fire together, wire together." Cells with high thresholds must also have high reinforcement values to overcome the initial thresholds. One can only do that through pooled inhibition by neighbors.

Lastly, cells in the visual cortex are not randomly assigned an orientation preference before the animal opens its eyes, instead there is some order to the layout. Some cells that prefer the same orientation will by clustered together, though these clusters will be randomly dispersed throughout the cortex. This variable is negatively correlated with the biological equivalent of the number of images the animal views. If these swaths of cells are large, then it will take more time steps to "overcome" their size, as only the cells on the edges of these swaths will be affected by the inhibition of their neighbors with different orientation preferences. As time continues, the edges of these swaths will become less and less uniform. However, the swaths do not completely disappear due to Threshold and reinforcement which are responsible for creating the stable iso-orientation domains.

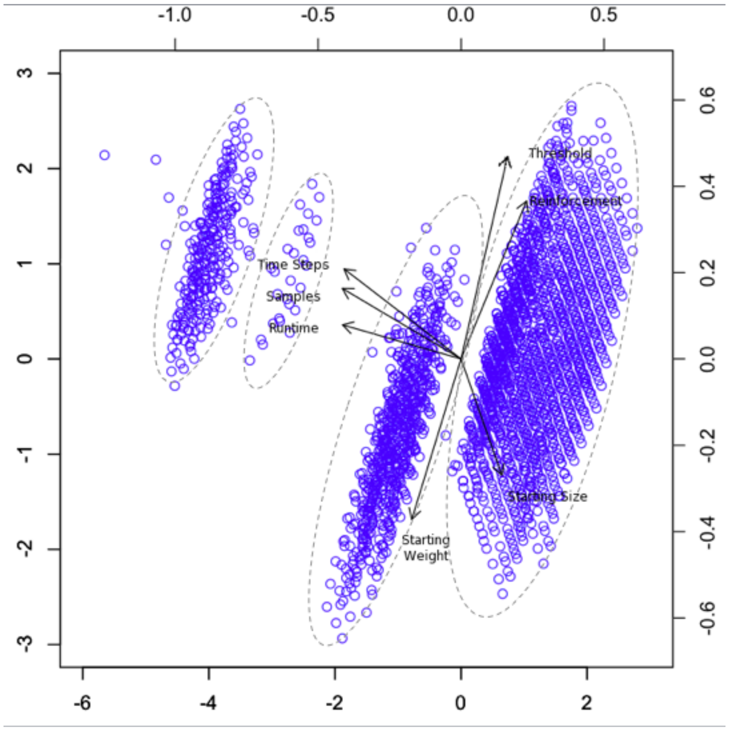

Figure 10. Four PCA clusters
The four clusters outlined.

As can be seen in Figure 10, the PCA yields four clusters of data. Two sets (the two to the left) are highly correlated with $\mathit{timesteps}$, $\mathit{samples}$, and $\mathit{runtime}$. Given more data, it is likely that these two clusters would form a single cluster as the range for number of time steps was filled out. These two clusters make up maps where $\mathit{startingsize}$ is negatively correlated with $\mathit{timesteps}$, $\mathit{samples}$, and $\mathit{runtime}$. Therefore, these maps likely have a small $\mathit{startingsize}$.

The third cluster (from the left) contains maps that are not so negatively correlated with $\mathit{startingsize}$. The fourth cluster contains maps which simply show the same negative correlation, but with the direction reversed: smaller time steps with a larger starting size. From this analysis, we can infer that with more time steps and/or samples, the influence of a larger starting size diminishes.

We also see that $\mathit{startingweight}$ is negatively correlated with $\mathit{threshold}$ and $\mathit{reinforcement}$. A large threshold or reinforcement value can diminish the influence of the starting weights applied to the cortical cells. $\mathit{Threshold}$ and $\mathit{reinforcement}$ are positively correlated: the larger the threshold to overcome, the larger the reinforcement must be to overcome it.

6.1.2 Evaluative Parameters

The complementary step to this analysis is to do a PCA on the evaluative measures, and then on this PCA, note which maps cluster near the real maps. Given this, we can match the name of the map to this set of input parameters, and see which input parameters create biologically similar maps.

Figure 13 shows the first two components of the PCA of the output measures for the two sets of data: data generated by the model (blue) and data generated from empirical maps (red). These measures include: Counterclockwise/Clockwise ratio, pinwheel CV, CW pinwheel CV, CCW pinwheel CV , and Centroid CV. Importantly, where the two sets of data overlap is where we find model input parameters that are capable of generating maps similar to real maps across all of these variables.

While the first two principal components capture most of the variance of the dataset, the other components do carry some variance that could differentiate clustered data from one another. Figure 12 shows the first three principal components. Unfortunately, a visual inspection suggests that the "real" data differs from the model data with the added third component. There may be some very partial overlap, but if there is any, it is on the fringe.

To explore whether or not this third component was important, a Scree Test was performed on the dataset. "Scree" is a term from geology referring to the rubble at the bottom of a cliff. We first take a correlation matrix, and decompose it into independent weighted combinations of the original variables (these combinations correspond to components). Each set will have some variance associated with it. The idea of the scree test is that if a factor is important, it

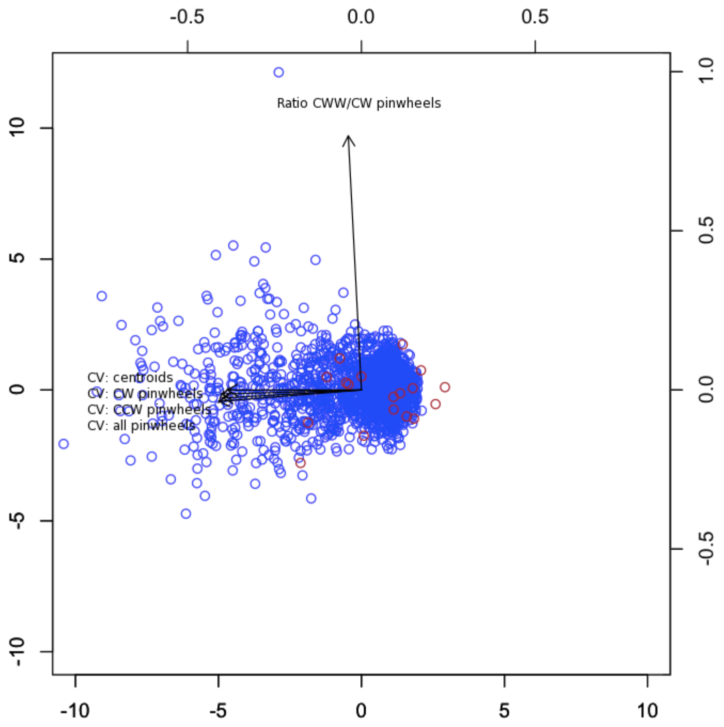

Figure 11. PCA of output measures
Data from model-generated maps are in blue and data from empirical maps are in red. Note the overlap of generated data to empirical data.

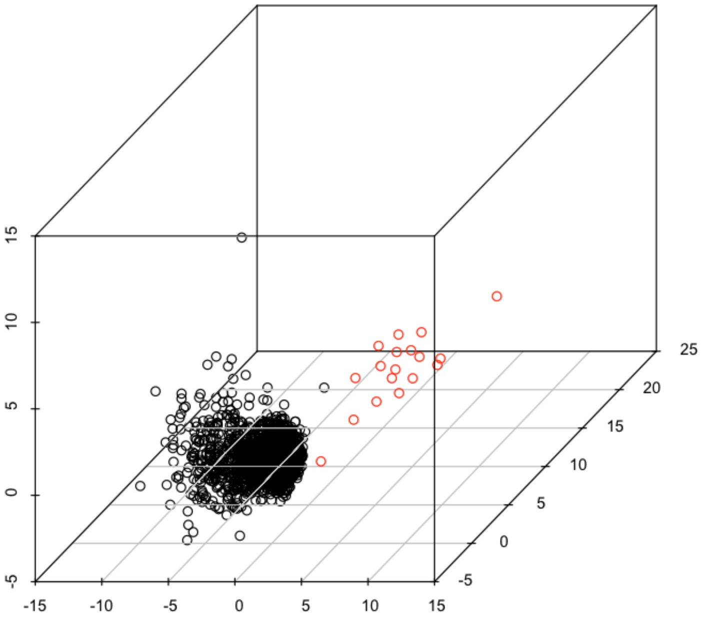

Figure 12. 3D PCA of output measures
Data from model-generated maps are in black and data from empirical maps are in red. Note that the overlap of generated data to empirical data is less when we add the third component.

will have a large variance. Therefore, we order the factors by variance, and plot the variance against the factor number. Then you keep the number of factors above the “elbow” in the plot. These are the important factors which account for the bulk of the correlations in the matrix, or the most salient principal components. It is called a scree test because the graph looks a bit like where a cliff meets the plain. In looking at a cliff, you might want to decide where the cliff stops and the plain begins. With the scree test you see where the important factors stop and the unimportant ones start. What we can do is to create a plot of the eigenvalues (variances) against their serial order (see **??**

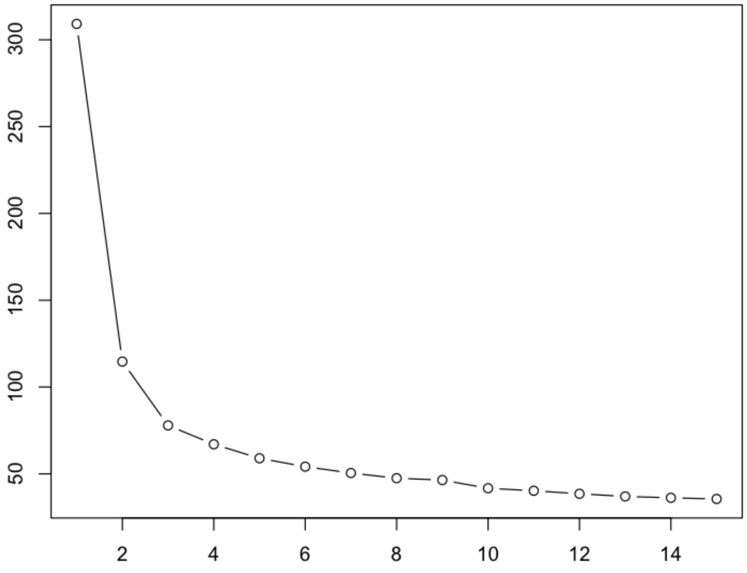

Figure 13. Scree Test
A Scree test on the data. The first two components contain a majority of the variance, and the “elbow” appears to be at around 3.

The Scree test suggests that there are 2 or possibly 3 underlying factors (meaning 2 or 3 principal components in the data are important), and that most of the variance is within the first two or three components. The slope between 2 and 3 is small, and so the decision may just be to conclude that there are two main factors involved. Based on this Scree Test, I will cluster the data according to the first two components using the K-means clustering algorithm.

K-means clustering is a method of cluster analysis which aims to partition n observations into k clusters in which each observation belongs to the cluster with the nearest mean. Figure 14 shows clustering of the data into two clusters. Over 82.5% of the in-between variance is accounted for when we cluster for two, meaning that most of the variance in the dataset is accounted for between the two (not within). This means they are fairly disjoint. Figure 15 shows clustering with three clusters. For three clusters, 68% of the variance is due to the difference between clusters. An analysis locating two or three clusters appear quite similar. These findings indicate that clusters 2 and 3 are likely quite similar, and I can present the data as having two main clusters.

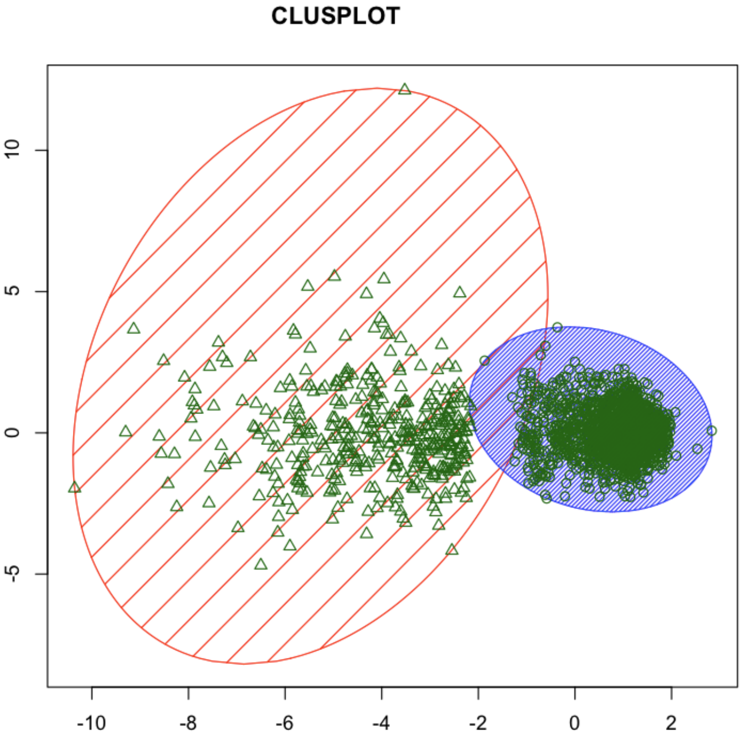

Figure 14. Two Cluster K-means
Clustering given two clusters

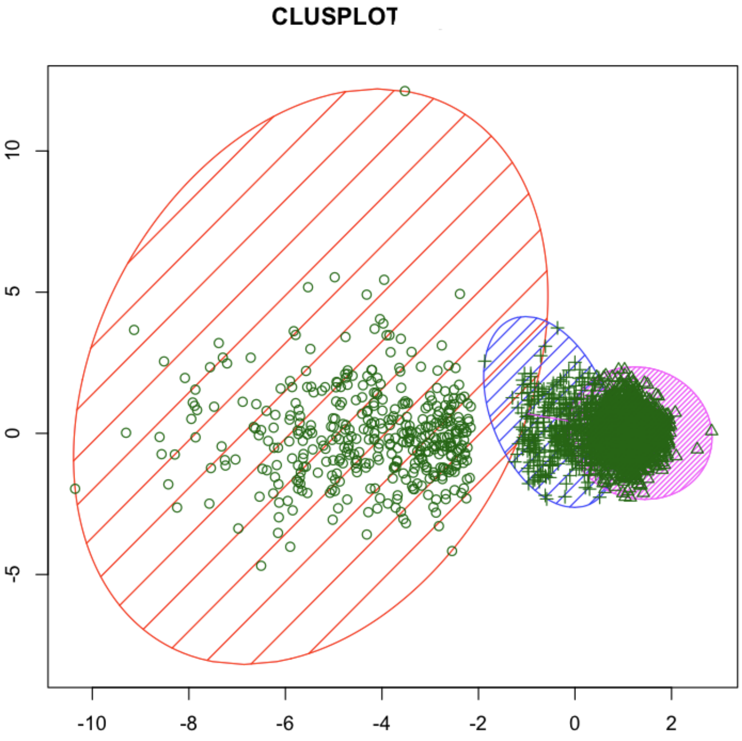

Figure 15. Three Cluster K-means
Clustering given three clusters

6.2 Run time and Accuracy

Figure 16 is a graph of runtime and accuracy.

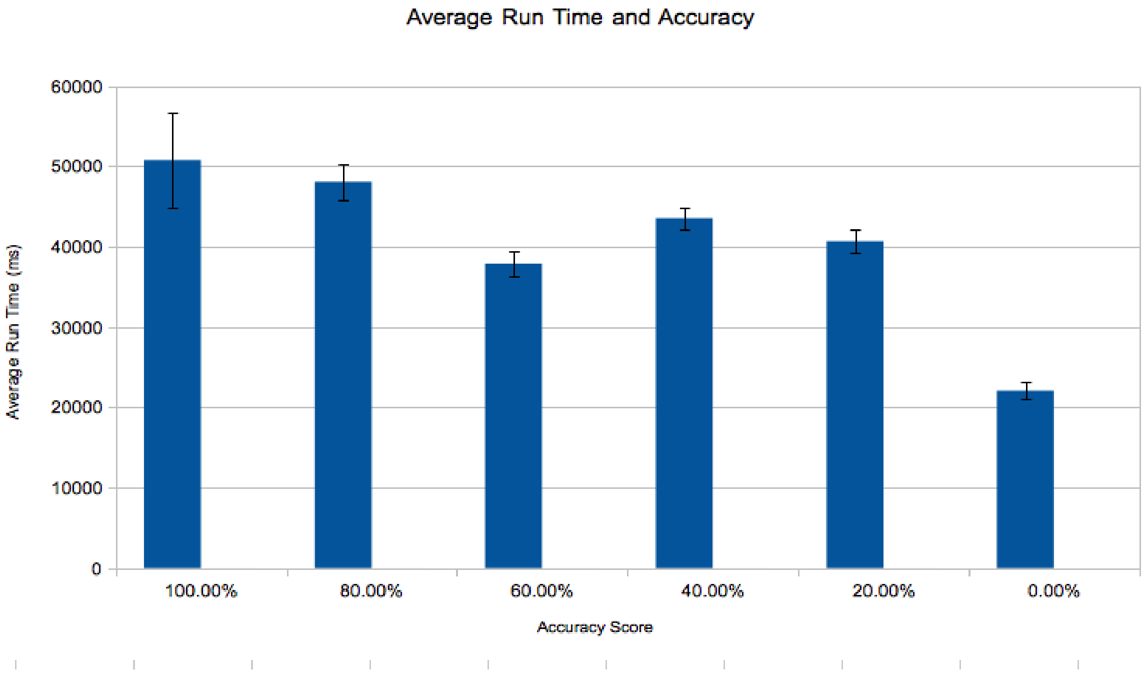

Figure 16. Accuracy Scores
Average run time of maps with the five different accuracy scores. Error bars represent 1 SEM. A longer run time tends to produce the most accurate maps.

There are five main measures in the map evaluation: the CCW/CW ratio, CV of pinwheels, CV of CCW pinwheels, CV of CW pinwheels, and the CV of the domain centroids. To determine a measure of "accuracy", the average of these five variables was taken on all of the real data.

If for the model map, a value for any one of these five variables fell within one standard error of the real mean (SEM), then one point (out of five) was awarded to the "accuracy total". Therefore, if any model map has an accuracy measure of "5" then that means the value for the variable fell within 1 SEM of the real map data average. The most accurate maps have an accuracy score of 5. This translates to an accuracy score of 100%.

The data have been summarized in Figure 16. On the x-axis is the accuracy score, and on the y-axis is the average run time (in ms) for all maps with the respective score. Please note that all time steps are equal, therefore run time might indicate the complexity of the computation and does not reflect the amount of time the model viewed the stimulus. The trend from the data suggest that longer run times produce more accurate maps. Further inquiry shows that the most accurate maps (those with 100% accuracy) take between 14 seconds to 1:37 minutes.

The data may dip at 60% due to a particular "sweet spot" range required for an accurate map, for some cases. For example, number of pinwheels plays a role in an accurate map. However, if the starting parameters were too large or too small, the number of pinwheels generated could increase with time. Perhaps with a larger run-time, we might see too many pinwheels, but only with particular parameters. Before attempting to publish this manuscript, I must address this possibility.

However, it must be mentioned that a fit of 100% may also not be accurate. This is called "over-fitting", meaning that I might be describing random error or noise instead of the underlying relationships I am looking for. Additionally, with this variance accounted for in my "accuracy" score, it is likely that I am over-fitting the data, and if I allowed some "noise" or did some more sophisticated significance testing, I may find a better way to describe the model's ability to generate "real" data. For example, by computing the accuracy score in this way, I am increasing the Familywise Error Rate, which is the probability of making one or more false discoveries, or type I errors among all the hypotheses when performing multiple hypotheses tests. If I treat each SEM range for each variable compared to another as a hypothesis test (testing whether or not they are significantly different), then in total I am performing 10 tests across all 5 variables. In short, with this method of comparison, a map may receive a high accuracy score due to chance.

Average RT	StDev RT	Min RT	Max RT	Count	Accuracy Score
50788	24840	14284	97749	18	100%
48058	24398	7007	107169	114	80%
37558	27403	6993	121715	341	60%
43558	33763	6991	311622	622	40%
40655	36094	6987	277809	679	20%
22098	27199	6983	152095	623	0%

TABLE I

TABLE OF RUN TIME AND ACCURACY STATISTICS

Table I summarizes the statistics for the run time and accuracy comparisons. This table shows that only 18 of the total 2397 maps achieved a "perfect" accuracy score, and that the minimum run time was nearly twice as what was required for the maps achieving 80% accuracy.

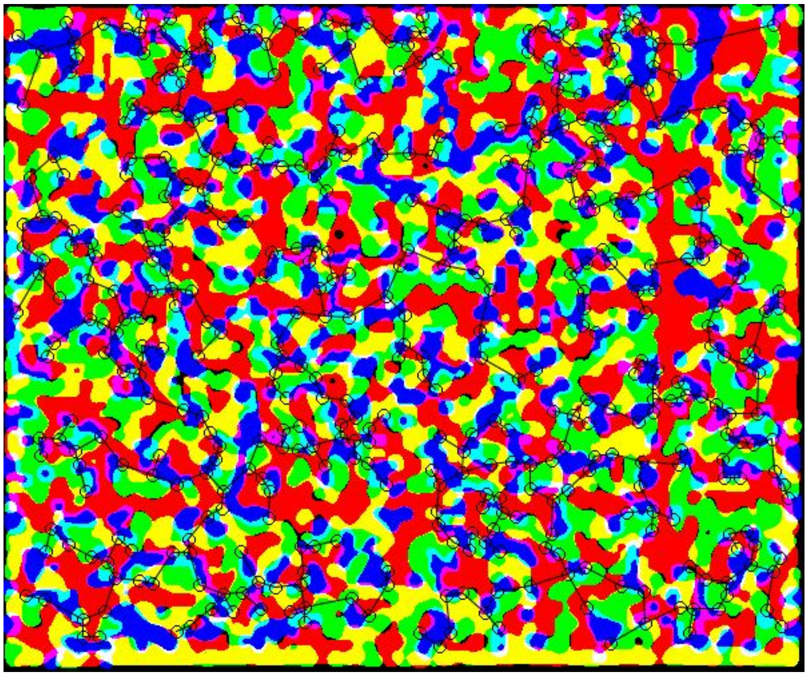

Figure 17. MST on a map generated by the model

Lastly, Figure 17 shows a model-generated map and Figure 18 shows an empirical map. Please note that the scales are completely different and only when compared in a normalized way can the two maps be considered similar. For example, the average compactness of Fig-

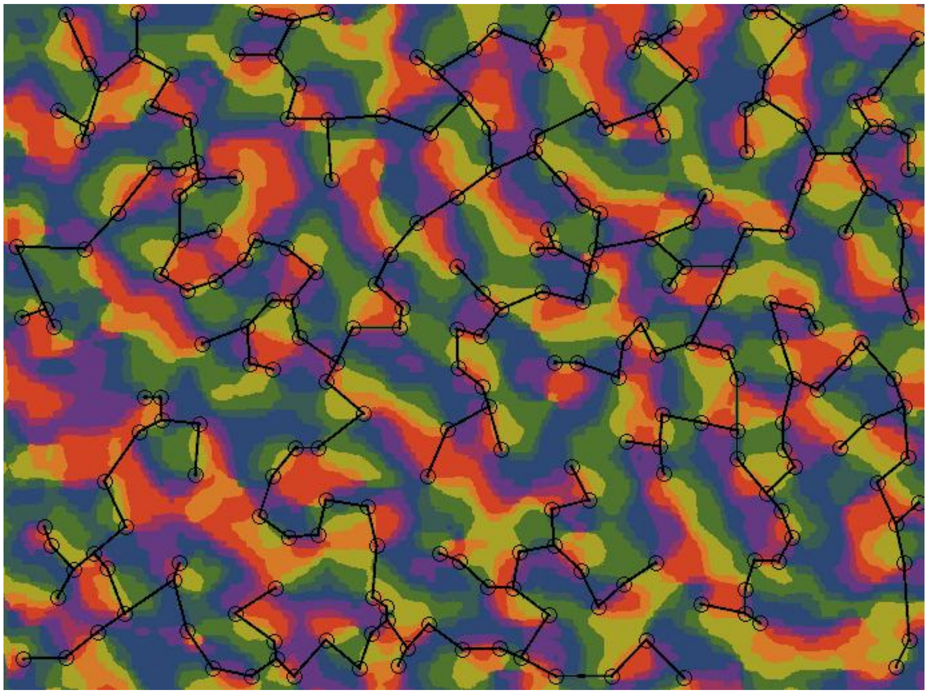

Figure 18. MST on a map imaged from monkey cortex

ure 17 is approximately 29, however the average compactness of Figure 18 is approximately 80. Therefore, the scale of Figure 17 to Figure 18 is about 3:8.

Given this analysis, the parameters to be input into this model to create the most biological similar maps are summarized in Table II. According to Table II, for this model, 10 time steps are sufficient to get biological similarity. This is a surprising result because I would expect that more images would further tune the map to be biologically similar. However, as few as 10 images seems to work well for this model. However, the mode is 20 time steps. Therefore, 20

Time Steps	Start Weight	Start Size	Reinforcement	Threshold
10	1	5	0	4
10	1	8	2	1
10	1	4	2	5
10	1	6	1	6
10	1	5	4	1
20	4	6	4	2
20	3	6	4	4
20	1	6	2	2
20	1	4	1	6
20	4	8	2	0
20	4	8	3	6
20	3	6	3	0
20	2	6	3	0
20	1	6	3	1
20	1	6	2	6
20	1	6	3	3
70	1	6	2	5

TABLE II

TABLE OF IDEAL PARAMETERS FOR THE 100% FIT images is the norm for this set of data. Please note that the data could also be ideal due to over-fitting or Type I errors.

A $\mathsf{StartingWeight}$ mode was 1 (ranging from 1 to 5), therefore it is important that the weighted orientation preferences of the cells at initialization are not too overpowering. The $\mathsf{StartSize}$ mode was 6 (ranging from 2 to 10), and so a starting size for the iso-orientation domain that is slightly over 5% the width of the retina width appears optimal. Reinforcement values (ranging between 1 and 7) seem to remain fairly low, therefore $\mathsf{reinforcement}$ cannot be overpowering either. Lastly, the $\mathsf{threshold}$ can vary, depending on the other parameters. In fact, there is a moderately negative correlation between $\mathsf{threshold}$ and $\mathsf{reinforcement}$ ($\mathsf{r} = -0.41$), and a positive correlation between $\mathsf{StartWeight}$ and $\mathsf{StartSize}$ ($\mathsf{r} = 0.51$).

CHAPTER 7

DISCUSSION

It is still unclear how the brain processes and stores information, and research into understanding “how” continues to evolve. The work presented here has demonstrated that by following the basic, organizational principles employed by the brain, one can achieve what would be considered biologically plausible responses.

7.1 Future Improvements

Future improvements for this model include adding additional dimensions of the visual experience such as direction of motion, color, spatial frequency, and direction selectivity. Additionally, while this model is composed of leaky-integrate and fire neurons, it should be tested with other types of neuron models. For example, there is another interesting representation of a neuron; one proposed by Bhaumik and Maher (14), which considers the resting and firing potential of the most recently fired neuron. This model allows for the integration of Long Term Potentiation (LTP) which is important for Hebbian learning adherence in the brain. The next iteration of the model should include this neuron:

$$u_I(t) = \eta(t - t_{fc}) + \beta_I \sum_j W_{ij}^{+/-} + \sum_{fL} \varepsilon(t - t_{fL}) + R_p$$

Where t_{fc} is the most recent firing time of the neuron, $\eta(t - t_{fc})$ is the refractory voltage of the cortical cell, $W_{ij}^{+/-}$ is the ON (+) or OFF (-) center cell weight between the LGN(j) and the cortical cell (i), $\varepsilon(t - t_{fL})$ is the EPSP (excitatory post-synaptic potential) in the cortical cell when the LGN cell fires, R_p is the resting potential, β_I is a function representing synaptic plasticity (the ability of the human brain to change as a result of one's experience), and where $u_I(t)$ is the membrane potential. The function for synaptic plasticity is as follows:

$$\beta_I = (H - L)f(A_v) + L$$

Where H and L are constants indicating the maximum and minimum values for the global scaling factor, A_v is the average input to the cortical cell I over time T.

$$A_v = \frac{1}{t}\int_o^t \sum_j W_{ij}X_j$$

$$f(A_v) = 1 - tanh[G(A_v - \theta)]$$

Here, G and θ are also constants and $X_j(t) \in 1, 0$. $X_j(t)$ is the firing of the LGN.

Additionally, improvements in the model can be made at the level of the dLGN. Currently, the model contains only neurons and are not differentiated by the system to which they belong (i.e. magnocellular or parvocellular). If in future iterations, I include sensitivity to color and luminance, it is important that I have a biological representation of these systems within the dLGN. Also, the dLGN in the current model contains only a single layer of neurons, but as

I include more dimensions of the visual experience, it is important I add these layers to the dLGN.

Lastly, at the level of the retina, I should incorporate both ON-center and OFF-center ganglion cells. Currently, I have only ON-center ganglion cells within the retina. While this is enough to discriminate the strong edges I present to the network, it will likely prove inferior to more natural scenes. Therefore, in the next iteration, having a set of OFF-center ganglion cells is important.

Most recently, Bednar et al (26) came out with what they are calling the most "biologically plausible" model of V1. It would be interesting to compare the accuracy of this model compared to mine. However, this model is different from mine in that theirs is considered the first developmental model with wiring consistent with V1 only, whereas my model considers signals from retina up to V1. My model is deficient in that it does not cover all feature dimensions covered in V1 as Bednar et al's does.

7.2 Conclusion

Despite the simplicity of this model, I have clearly demonstrated that following very simplistic rules of basic neuron firing and connecting them in a biologically relevant way, a map of orientation-selective cells akin, to what would be seen in nature, can be generated. This goes to show that when considering models of the brain, it is very important to take into account the architecture and subsequent wiring of afferent and efferent signals into different areas when developing models of the brain.

I have found what parameters are positively and negatively correlated. From this example, we can take away that neuron units in the brain, when placed within a network, must find a particular balance between firing threshold and reinforcement value. In the biological brain, this threshold is determined by ionic receptors on the cell membrane. The more receptors on the membrane than the faster the cell may reach threshold. If a cell is quick to reach threshold, then it is more apt to fire, so perhaps the reinforcement value when it does fire is greater. This effect in a real neural network is known as LTP. This is an unexpected but very interesting finding. LTP is an emergent property of this model.

Growth in the area of computer architecture is booming, and companies that seek to model the brain are further relying on redefining current computer architecture and its underlying processes. They are finding that it is important to keep processes parallel, yet to have these processes also within a very specific hierarchy. With these components being built to mimic the firing of biological cells such as neurons, computational power is increasing such that the reality of building a computational "brain" may be on the horizon.

CITED LITERATURE

1. Neumann, J. V.: The Computer and the Brain. Yale University Press, 2012.

2. Levine, M.: Sensation and Perception. Oxford University Press, 2000.

3. Cisek, P., Drew, T., and J.F. Kalaska, E.: Computational Neuroscience: Theoretical Insights into Brain Function. Elsevier, 2007.

4. Holland, J.: Emergence: From Chaos to Order. Oxford University Press, 1998.

5. Kandel, E., Schwartz, J., and Jessell, T.: Principles of Neuroscience, 4th Ed.. McGraw-Hill, 2000.

6. Ho, K., Roessmann, U., Straumfjord, J., and Monroe, G.: Analysis of brain weight. adult brain weight in relation to sex, race, and age. Archives of Pathology and Laboratory Medicine, 12:635–639, 1980.

7. Markram, H.: The blue brain project. Nature Reviews Neuroscience, 7:153–160, 2006.

8. Campbell, M., Jr., A. H., and Hsu, F.: Deep blue. Artificial Intelligence, 134:57–83, 2002.

9. Hubel, D. and Wiesel, T.: Functional architecture of macaque monkey visual cortex. Proc R Soc Lond B, 198:1–59, 1977.

10. Blasdel, G. and Salama, G.: Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. Nature, 321:579–585, 1986.

11. Braitenberg, V. and Braitenberg, C.: Geometry of orientation columns in the visual cortex. Biological Cybernetics, 33:179–186, 1979.

12. Bonhoeffer, T. and Grinvald, A.: Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. Nature, 353:429–431, 1991.

13. Gotz, K.: Do “d-blo” and “l-blob” hypercolumns tessellate the monkey visual cortex? Biological Cybernetics, 56:107–109, 1987.

14. Bhaumik, B. and Mathur, M.: A cooperation and competition based simple cell receptive field model and study of feed-forward linear and nonlinear contributions to orientation selectivity. Journal of Computational Neuroscience, 2003.

15. Schummers, J., Sharma, J., and Sur, M.: Bottom-up and top-down dynamics in visual cortex. Progress in Brain Research, 149:65–81, 205.

16. Hubel, D. H. and Wiesel, T. N.: Receptive fields and functional architecture of monkey striate cortex. Journal of Physiology, pages 215–243, 1968.

17. Bhaumik, B., Agarwal, A., and Manohar, M.: Receptive field properties of near neighbor orientation selective neurons in the visual cortex: a modeling study. International Journal of Neural Systems, 15(1 & 2):31–40, 2005.

18. Valois, R. D. and Valois, K. D.: Spatial Vision. Oxford University Press, 1990.

19. Erwin, E., Obermayer, K., and Schulten, K.: Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. Neural Computation, 7:425 – 468, 1995.

20. Swindale, N. and Bauer, H.: Application of kohonen's self-organizing feature map algorithm to cortical maps of orientation and direction preference. In Proceedings of the Royal Society London Biology, volume 265, pages 827–838, 1998.

21. Swindale, N.: How different feature spaces may be represented in cortical maps. Network, 15:217–242, 2004.

22. McCulloch, W. and Pitts, W.: A logical calculus of the ideas immanent in nervous activity. Bulletin of Mathematical Biophysics, 5:115–133, 1943.

23. Wiesel, T. N. and Hubel, D.: Single-cell responses in striate cortex of kittens deprived of vision in one eye. Journal of Neurophysiology, 26:1003–1017, 1963.

24. Hubel, D. and Wiesel, T. N.: Binocular interaction in striate cortex of kittens reared with artificial squint. Journal of Neurophysiology, 28:1041–1059, 1965.

25. Wiesel, T. N. and Hubel, D.: Extent of recovery from the effects of visual deprivation in kittens. Journal of Neurophysiology, 28:1060–1072, 1965.

26. Bednar, J.: Building a mechanistic model of the development and function of the primary visual cortex. Journal of Physiology(Paris), 106:194–211, 2012.

VITA

NAME	Jennifer E. Anderson
EDUCATION	B. A., Psychology, University of Illinois at Chicago, Chicago, Illinois, 2007 M. A., Behavioral Neuroscience, University of Illinois Department of Psychology, Chicago, Illinois, 2010
TEACHING	Department of Psychology, University of Illinois, Chicago, Illinois Department of Computer Science, University of Illinois, Chicago, Illinois
PROFESSIONAL MEMBERSHIP:	AANS Science Vision Sciences Society
ABSTRACTS	**Anderson J**, Anand A, Franzini S, Levine M, and Wirtschafter, D. (2010) A means of map measurement: orientation specificity in primary visual cortex. SfN Poster Session: Poster 108 **Anderson J**, and Levine M. (2010) Noise Modulation in the Dorsal and Ventral Visual Pathways [Abstract]. VSS Poster Session: *Journal Of Vision*, 10(7): 1075. Anand, A, **Anderson J**, Berger-Wolf, T. (2010) Predicting Orientation Selectivity in Primary Visual Cortex [Abstract]. VSS Poster Session: *Journal Of Vision*, 10(7): 936. **Anderson J**, Levine M, McAnany, JJ. (2009) Response demands do not influence perceived illusory motion in cognitive-based tasks [Abstract]. VSS Poster Session: *Journal Of Vision*, 9(8): 842, 842a

ABSTRACTS CON'T

Levine M, **Anderson J**, McAnany, JJ. (2009) Two modes of hiding suprathreshold stimuli in complex patterns [Abstract]. VSS Poster Session: *Journal Of Vision*, 9(8): 1006, 1006a.

Anderson, J, Levine, M, McAnany, JJ. (2008) Prestidigitation: Easier to fool the eye than the hand [Abstract]. VSS Poster Session: *Journal Of Vision*, 8(6): 1051, 1051a.

Levine, M, McAnany, JJ, **Anderson, J**. (2008) The effect of curvature on the grid illusions: Influence of a homunculus? [Abstract]. VSS Poster Session: *Journal Of Vision*, 8(6): 278, 278a.

PUBLICATIONS

Levine, ML, **Anderson, JE,** and McAnany, JJ. (2012) Effects of orientation and contrast upon targets in straight and curved arrays. *Perception*, 41, 1419-1433.

Veilleux, J.C., Colvin, P.J., **Anderson, J.**, York, C., & Heinz, A.J. (2009). A review of opioid dependence treatment: Pharmacological and psychosocial interventions to treat opioid addiction. Clinical Psychology Review.