

**Modeling the self-organization
of color-selective neurons in
the visual cortex**

Judah Ben De Paula

Report AI-TR-07-347 August 2007

judah@cs.utexas.edu
<http://www.cs.utexas.edu/users/mn/>

Artificial Intelligence Laboratory
The University of Texas at Austin
Austin, TX 78712

Copyright

by

Judah Ben De Paula

2007

The Dissertation Committee for Judah Ben De Paula
certifies that this is the approved version of the following dissertation:

**Modeling the self-organization of color selectivity
in the visual cortex**

Committee:

Risto Miikkulainen, Supervisor

James A. Bednar, Supervisor

Benjamin Kuipers

Bruce Porter

Eyal Seidemann

**Modeling the self-organization of color selectivity
in the visual cortex**

by

Judah Ben De Paula, B.S.

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

August 2007

In memory of Bruce Lyle De Paula, J.D., who knew the answer when I asked,
“Daddy, why is the sky blue?”

Acknowledgments

I would like to thank my advisor Risto Miikkulainen for his continuous interest, support, and encouragement throughout this research. An hour talking with Risto, and I felt like the anything could be accomplished.

A special thank you must go to James Bednar, my committee co-chair, who has been with me every step of the way. Jim was the one who first introduced me to LISSOM, then guided my exploration, and provided irreplaceable and critical technical advice throughout the development of my dissertation. My thesis builds directly upon the foundation he laid, and I could not have accomplished what I did without his aid.

I am also thankful for the interest and support of the committee members Benjamin Kuipers, Bruce Porter, and Eyal Seidemann. Eyal Seidemann in particular has helped keep this computer scientist honest as I wrote about the brain.

Thank you to Yiu Fai Sit, who walks the same path I do, and has helped with the directions. I wish him the best in reaching his own finish-line. Thank you to Lisa Kaczmarczyk for her friendship, advice, understanding, and the washer and dryer she left me when she graduated. Thank you to Erik Reeber who for four years was the best roommate a guy could have, that is, until I met my wife.

Thank you to my wife, Brooke, who has helped to keep my subjects and verbs in agreement.

I am grateful to all my colleagues in the UTCS Neural Networks group, both past and present, who enriched my academic career with their conversation and friendship. I appreciate their feedback to my many talks given in the Neural Network group meetings which has made me a better presenter.

JUDAH BEN DE PAULA

The University of Texas at Austin
August 2007

Modeling the self-organization of color selectivity in the visual cortex

Publication No. _____

Judah Ben De Paula, Ph.D.

The University of Texas at Austin, 2007

Supervisor: Risto Miikkulainen
and James A. Bednar

How does the visual cortex represent and process color? Experimental evidence from macaque monkey suggests that cells selective for color are organized into small, spatially separated blobs in V1, and stripes in V2. This organization is strikingly different from that of orientation and ocular dominance maps, which consist of large, spatially contiguous patterns. In this dissertation, a self-organizing model of the early visual cortex is constructed using natural color image input. The modeled V1 develops realistic color-selective receptive fields, ocular dominance stripes, orientation maps, and color-selective regions, while the modeled V2 also creates realistic color-selective and orientation-selective neurons. V1 color-selective regions are generally located in the center of ocular dominance stripes as they are in biological maps; the model predicts that color-selective regions become more widespread in both cortical regions when the amount of color in the training images is increased. The model also predicts that in V1 there are three types of color-selective regions (red-selective, green-selective, and blue-selective), and that a unique cortical activation pattern exists for each of the HSV colors. In both V1 and V2, when regions of different color-selectivity are located nearby, bands of color form with gradually changing color preferences. The model also develops lateral connections between cells that are selective for similar

orientations, matching previous experimental results, and predicts that cells selective for color primarily connect to other cells with similar chromatic preferences. Thus the model replicates the known data on the organization of color preferences in V1 and V2, provides a detailed explanation for how this structure develops and functions, and leads to concrete predictions to test in future experiments.

Contents

Acknowledgments	vi
Abstract	vii
List of Tables	xiii
List of Figures	xiv
List of Equations	xvi
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Approach	3
1.3 Results	4
1.4 Outline of the dissertation	5
Chapter 2 Background	7
2.1 Color perception	7
2.2 The visual pathway	8
2.2.1 Photoreceptors in the retina	9
2.2.2 Retinal Ganglia and Lateral Geniculate Nucleus	11
2.2.3 Primary visual cortex	14
2.2.4 Secondary visual cortex	18
2.3 Developmental input-driven self-organization	19
2.3.1 Visual cortex studies	19
2.3.2 Synaptic plasticity and the Hebb principle	20
2.3.3 Homeostatic plasticity	20
2.4 Self-organizing models of the visual cortex	21
2.4.1 Monochrome vision models	21
2.4.2 Models of color selectivity	22
2.5 Conclusion	23

Chapter 3	The LISSOM model	24
3.1	Architecture	24
3.1.1	Eyes	26
3.1.2	Lateral geniculate nucleus	26
3.1.3	Primary visual cortex	28
3.1.4	Secondary visual cortex	28
3.2	Running the model	28
3.2.1	Activation	30
3.2.2	Cortical settling	30
3.2.3	Learning	31
3.3	Training stimuli	31
3.4	Conclusion	32
Chapter 4	Statistics of natural images	33
4.1	Correlations between photoreceptor types	34
4.2	Independent and shared information between photoreceptor types	35
4.2.1	Joint entropy	35
4.2.2	Mutual information	38
4.2.3	Minimum unique entropy	39
4.2.4	Redundancy	39
4.2.5	Discussion	41
4.3	Color distribution	41
4.4	Principal component of image patches	42
4.4.1	The simulation architecture	44
4.4.2	Results	45
4.4.3	Discussion	45
4.5	Color and luminance gradients	48
4.6	Conclusion	50
Chapter 5	Modeling self-organizing feature preferences in primary visual cortex	52
5.1	Map level measurements	52
5.1.1	Orientation-selectivity maps	53
5.1.2	Color maps	53
5.1.3	Ocular dominance	55
5.2	Individual neurons	57
5.2.1	Color-selective neurons	57
5.2.2	Orientation-selective neurons	59
5.2.3	Eye-selective neurons	61
5.3	Discussion	62

5.4	Conclusion	63
Chapter 6 Modeling self-organizing feature preferences in secondary visual cortex		64
6.1	Map level measurements	64
6.1.1	Orientation maps	65
6.1.2	Color maps	65
6.1.3	Ocular dominance maps	66
6.2	Individual neurons	69
6.2.1	Color-selective neurons	69
6.2.2	Orientation-selective neurons	73
6.3	Discussion	73
6.4	Conclusion	74
Chapter 7 Experimental Analysis		75
7.1	Sensitivity to image features	75
7.1.1	Channel correlation	76
7.1.2	Feature map interaction	76
7.2	Random color distributions	77
7.2.1	Training stimuli	77
7.2.2	Results	78
7.3	Self-organizing the full color spectrum	81
7.3.1	Training stimuli	81
7.3.2	Results	82
7.4	Discussion	82
7.5	Conclusion	84
Chapter 8 Discussion and Future Research		86
8.1	Contributions of the dissertation	86
8.2	The trichromatic LISSOM model	87
8.2.1	LISSOM architecture	87
8.2.2	Future work	88
8.3	Training stimuli and natural images	91
8.3.1	Colors in the World	91
8.3.2	Future work	92
8.4	Primary visual cortex	93
8.4.1	Predictions of the model	93
8.4.2	Future work	94
8.5	Secondary visual cortex	96
8.5.1	Inputs and color wheels	96

8.5.2	Predictions of the model	96
8.5.3	Future work	97
8.6	Conclusion	97
Chapter 9	Conclusion	99
9.1	Summary	99
9.2	Conclusion	101
Appendix A	Converting RGB images to LMS cone activations	102
A.1	RGB to LMS algorithm	102
A.2	Python conversion code	105
Appendix B	Simulation parameter values	106
B.1	Overview	106
B.2	V1 and V2 default parameters	107
B.3	New retina and V2 parameters	108
B.4	V1 and V2 parameter schedules	108
B.5	Unique V2 simulation parameter values	111
B.6	Extending a monochrome LISSOM model to color	114
Bibliography		115
Vita		130

List of Tables

4.1	Correlations between channels in various image corpora	35
4.2	Joint entropy of channels in various image corpora	37
4.3	Mutual information of channels in various image corpora	37
4.4	Minimum unique entropy of channels in various image corpora	38
4.5	Channel redundancy in various image corpora	41
B.1	Parameters for the LISSOM simulation	107
B.2	Defaults for V1 and V2 constant parameters	110
B.3	New retina and V2 parameters for color simulations	111
B.4	V1 and V2 parameter modification schedule	112
B.5	V2 region parameter initialization	113

List of Figures

1.1	California hillside 1: Monochrome	2
1.2	California hillside 2: Color	5
2.1	Color perception: Wavelengths and wheels	8
2.2	The visual pathway	9
2.3	Wavelength sensitivity functions	10
2.4	Center-surround receptive field	12
2.5	Red-green ON and OFF center-surround receptive fields	13
2.6	V1 feature preference maps	15
2.7	Lateral connections of tree shrew V1 neurons	17
2.8	V2 schematic organization	18
3.1	LISSOM model of orientation, long/medium/short color selectivity, and ocular dominance	25
3.2	Sample receptive fields for units in LGN sheets in the LISSOM model	27
3.3	LISSOM model	29
4.1	Graphical comparison of channel correlations in Table 4.1	36
4.2	Minimum unique entropy of RGB and LMS natural image sets	40
4.3	Hue histograms of natural scenes	43
4.4	Sample training input patches	45
4.5	Hebbian neuron MATLAB code.	46
4.6	Learned weight function of the Hebbian neuron	47
4.7	Gradient maps of a single natural image	49
4.8	Average gradients of image corpora	50
5.1	Self-organized orientation and color maps	54
5.2	Pure hue color-selectivity	55
5.3	Self-organized color maps and ocular dominance	56
5.4	Comparing biological and simulated V1	57
5.5	Receptive field of a color-selective neuron (red)	58

5.6	Receptive field of an orientation-selective neuron	60
5.7	Receptive field of an eye-selective neuron (blue color, right eye) . . .	61
6.1	V2 cortex trained using natural color images	66
6.2	V2 hue preference and selectivity maps	67
6.3	V2 eye preference and selectivity maps	68
6.4	Afferent receptive fields of V2 neurons encoded for hue preference . .	70
6.5	Lateral connections of three V2 color-selective neurons	71
6.6	Afferent receptive fields of V2 neurons encoded for orientation preference	72
7.1	Natural image training stimuli with rotated hues	79
7.2	Color-selectivity maps of V1 trained with rotated hues	80
7.3	V1 hue-selectivity map self-organized using artificial-color images . .	83
7.4	V2 color maps trained with artificial-color images	84
8.1	Retinal cone geometry	89

List of Equations

2.1	Hebb learning rule	20
3.1	LISSOM Unit Activation	30
3.2	LISSOM Cortical Unit Settling	30
3.3	LISSOM Cortical Learning	31
4.1	Joint entropy	35
4.2	Mutual information	38
4.3	Minimum unique entropy	39
4.4	Normalized redundancy	40
4.5	Hebbian connection weight normalization	44
4.6	Sum of all input patch pixels	46
4.7	Sum of all red input patch pixels	46
4.8	Sum of all green input patch pixels	47
4.9	Sum of all blue input patch pixels	47
4.10	Euclidean contour gradient distance measure	48
4.11	Hue contour gradient distance measure	48
A.1	A pixel's largest wavelength energy	104
A.2	Normalization of phosphor emission vectors	104
A.3	Total luminance power spectrum of one RGB pixel	104
A.4	Cone normalization: Independent channel modulation	105
A.5	Cone activation from viewing RGB pixel	105

Chapter 1

Introduction

Color is a vital and integral part of human experience. For instance, color vision allows us to effortlessly tell the difference between a ripe yellow banana and an unripe green one. Without color vision we could still tell the difference by taking a bite, but seeing green prevents an unpleasant experience. Or, imagine that we need to collect a medicinal plant that has a purple flower. The color of the flower will help us quickly identify the plant in a large field. To illustrate, Figures 1.1 and 1.2 show an unimpressive hillside that reveals an impressive variety of foliage—but only when seen in color.

1.1 Motivation

Although color is a physical property of the environment, it is perceived through a complex neural and cognitive process. The mind filters, processes, and organizes the information in the world so that we see color. This dissertation is a study of how the brain uses the activity pattern of neurons in the eye to develop color-selective cortical neurons that the higher regions of the brain can then use for color perception.

Most computational vision simulations to date have been monochromatic, i.e., they have not included color information. Lower animals do not have color vision, so a monochrome simulation is an appropriate model of these animals. However, in higher animals such as humans there are a large number of color-sensitive cells, and they convey important information about the visual environment. In order to build artificial vision systems with human-like processing of color stimuli it is necessary to understand how the brain represents and processes color.

Optical imaging of cortical activity in mammals shows that there are overlapping retinotopic feature maps in the primary visual cortex (Blasdel and Campbell, 2001; Issa, Trepel, and Stryker, 2001; Tootell, Silverman, and De Valois, 1981; Tootell,



Figure 1.1: **California hillside 1: Monochrome.** This figure is a black and white photograph of a hillside near Bakersfield, California. A monochromatic world provides nearly all necessary information to interact with the environment, such as slopes of the hills and textures in the vegetation. But without color, vital features of the world are not apparent. This image is in color in Figure 1.2 on page 5. Before turning the page, try to estimate how many different types of vegetation are growing here. (Photograph by Kee, 2005; reprinted with permission.)

Silverman, Hamilton, Switkes, and De Valois, 1988; Vanduffel, Tootell, Schoups, and Orban, 2002). That is, the brain displays systematic and spatially organized neuronal activity when the retina is presented with stimuli such as colors and bars of light. Feature maps are constructed by recording the strength of activity in the surface of the cortex in response to these different features. Each of the functional maps show how cortex contains distinctive cortical patterns for visual features such as orientation, ocular dominance, or spatial frequency. However, color appears to be organized differently from orientation and ocular dominance maps. Color-selective regions form blob-like patches co-located with cortical regions that stain for cytochrome oxidase (Hubener and Bolz, 1992; Landisman and Ts'o, 2002a,b; Sato, Katsuyama, Tamura, Hata, and Tsumoto, 1994). Color blobs are primarily found within V1 ocular dominance stripes (Bartfeld and Grinvald, 1992). While the maps for ocular dominance and orientation selectivity are systematic and continuous (i.e. hierarchically organized

patches and stripes), color maps consist of scattered clusters of selectivity, surrounded by large regions not selective for color. Why is the functional structure for color selectivity so different from the other features? Is there a computational or statistical reason why it would be mapped differently?

It has been difficult to obtain a complete understanding of color processing based on biological experiments. In part this is because it is impractical to characterize every neuron in the visual system precisely, because the numbers of neurons and connections are many orders of magnitude larger than can be measured directly. In addition, many of the experimental techniques required are invasive, and thus destroy the natural behavior of the experimental subject.

Computational simulations make it possible to test color vision theories and also provide predictions for future experiments. Simulations are an important complement to experimental approaches: all parameters of the system are known, the input stimuli can be reproduced exactly, and measurements can be made without disturbing the system itself. Computational simulation is therefore the approach taken in this dissertation.

1.2 Approach

This dissertation presents a study of a computational model of the early visual cortex that self-organizes to form color-selective regions. The underlying theory is that color selectivity is driven by activity patterns in the retina. This model, LISSOM (Laterally Interconnected Synergetically Self-Organizing Model; Miikkulainen et al., 2005), is a hierarchy of two-dimensional neural networks with Hebbian modification of afferent and lateral connections. Previous LISSOM simulations have successfully modeled how orientation selectivity and ocular dominance develop (Bednar and Miikkulainen, 2006; Miikkulainen, Bednar, Choe, and Sirosh, 2005), but color selectivity has not been included until now.

The color LISSOM model is trichromatic, includes both the primary and secondary visual cortex, and is validated with data from optical imaging experiments in the macaque monkey. Macaques have trichromatic vision similar to that of humans, and data is available about the spatial relationships among their orientation, ocular dominance, and color maps. That is, the structural relationships among the maps are known, even though it is unclear why they are each organized differently. By comparing the statistics of the training stimuli to the final self-organization of the cortex, the model can be used to explain why feature maps organize in different ways.

Before network training, the model cortex contains neurons with unorganized connection weights, and are unselective to features in the input stimuli. As each successive training stimulus is presented to the model retina, the neurons adapt their

connection weights and begin to form selectivity for features in the inputs. At the end of the training session, the simulated cortex has many of the same structures and features as in the brain.

1.3 Results

The trichromatic LISSOM model creates color-selectivity maps, ocular dominance maps, and orientation maps that are biologically realistic. Both in the brain and in the model, color-selective regions form patchy blob-like structures, and orientation maps form pinwheels, fractures, and linear zones. In addition, the interactions among the feature maps are also similar to known physiology. For example, the model shows color-selective regions interspersed between orientation-selective regions and within the ocular dominance stripes. As in biology, individual neurons develop preferences for input features such as oriented bars of light, or color (Hubel and Livingstone, 1983). Interestingly, the developed neuron preferences are not limited to colors at the photoreceptor sensitivity peaks, but also include the perceptually created magenta hues.

In addition to giving a computational explanation for biological observations, the model makes concrete predictions for future experiments. First, it suggests that proper cortical development requires that the orientation and color features in the input must be properly balanced in relation to the left and right eye input correlation. That is, luminosity gradients need to be a stronger feature of the stimuli than color gradients, and the inputs must have a strong correlation between the eyes.

Second, the model suggests a mechanism by which neurons can become selective for hues such as yellow for which there are no photoreceptors. Yellow-selective neurons have been identified experimentally, but the mechanism that gives rise to such preferences, given the output of retinal ganglion cells, is still unknown. Corresponding to retinal ganglion cells, color-selective neurons in the model V1 can be labeled red, green, or blue. Each color-selective blob primarily contains a single type of color-selective neuron. However, neurons in areas where red and green-selective blobs meet respond well to both red and green stimuli. As a result, they become selective for yellow, a hue that activates both red and green photoreceptors. A similar mechanism may take place in biology, giving rise to yellow-selective cortical neurons.

Third, the model predicts that neurons that are selective to similar colors will have strong lateral connections. These connections between them reflect long-term patterns of correlation present during self-organization, and act to create sparse activity patterns and smooth feature maps.

This dissertation makes three major contributions. First, the computational model helps explain how a visual hierarchy emerges in the brain, and how it con-



Figure 1.2: **California hillside 2: Color.** This is the original color photograph of the California hillside shown in black and white on page 2 in Figure 1.1. Color provides clear separation of the multiple types of vegetation. Yellow, orange, purple and red flowering vegetation can now be distinctly identified. Non-flowering vegetation can also be categorized more easily with deep green grasses separated from gray-green brush. Color provides unique information about the world that either cannot be gained from other visual features, or can only be obtained with greater difficulty. (Photograph by Kee, 2005; reprinted with permission.)

tributes to color processing. Second, computer science is given a biologically inspired color vision architecture that is both robust and scalable. Third, the descriptive power of the model provides a framework to discuss the visual system and enables additional modeling of how biological mechanisms interact. Therefore, a computational model of color vision expands our scientific knowledge and helps us understand the physiological basis of color perception.

1.4 Outline of the dissertation

This dissertation is organized into four parts: Background (**Chapters 1 and 2**), Model (**Chapter 3**), Results (**Chapters 4 through 7**), and Discussion (**Chapters 8 and 9**).

Chapter 2 surveys the vision pathway starting at the retina and moving through the lateral geniculate nucleus and primary visual cortex to the secondary visual cortex. Computational models are also reviewed with an emphasis on self-organizing cortical map simulations and how they have been applied to color-selective map modeling.

Chapter 3 presents the LISSOM model architecture and the necessary extensions to the retina and lateral geniculate nucleus to make a computational model that has realistic trichromatic photoreceptor activation.

Chapter 4 discusses five different techniques used to measure the latent statistics contained within natural images from around the world and used as network training stimuli.

Chapter 5 shows the self-organized primary visual cortex of LISSOM neurons trained with natural images. The color, orientation, and ocular dominance feature maps all have characteristics of biological neurons. In addition, the model makes predictions about biology that can now be tested in animals.

Chapter 6 presents results from a secondary visual cortex region that is self-organized using the LISSOM primary visual cortex as input. Color-selective and orientation-selective neurons develop that have similarities with biological neurons. The LISSOM model also makes predictions about the color-selectivity and lateral connectivity of V2 neurons.

Chapter 7 reveals that the primary and secondary visual cortex maps will develop additional color preferences if the natural images used as input are modified to have different color statistics. These control simulations prove that the LISSOM model is capable of preference maps that have all hues represented, and suggest that the cortex may have additional homeostatic regulatory mechanisms that achieve a desired balance among the colors in natural stimuli.

Chapter 8 discusses the predictions of the model and how they can be tested through biological experiments. Further, future work is discussed about how the model will be used to address additional biological questions, and applied to robotic vision tasks.

Chapter 9 concludes the dissertation with a summary of the results covered in the preceding chapters and the contributions made by the thesis.

Appendix A presents the algorithm to convert RGB images into simulated long, medium, and short photoreceptor activations. This transformation is used by the model to convert bitmap images into retinal photoreceptor activations.

Appendix B contains tables of the LISSOM parameter values used in the self-organizing simulations making it possible to replicate the work exactly.

Chapter 2

Background

The model presented in the following chapters simulates the structures found in the early visual system of trichromatic primates. This chapter contains background information useful for understanding the model architecture and methodology. First, the perceptual organization of color is presented. The second section then presents the physiological organization of the retina, lateral geniculate nucleus, and early visual cortex, each modeled in this dissertation. Third, the biological support for an input-driven model of cortical self-organization is presented. Finally, the fourth section surveys other computational models of visual cortex maps, discussing how they relate to the one in this thesis.

2.1 Color perception

In preparation for presenting results showing perceptually ordered color-selective neurons, the structure of the visible wavelength spectrum and the color wheel will be reviewed.

The visible light wavelength spectrum has red light at the long-wavelength end of the spectrum at about 650nm (Figure 2.1a), while at the other end of the visible spectrum is violet (Itten, 1970). A circle results from linking the ends of the color spectrum with magenta. Curiously, the perceptual color magenta does not exist as a single wavelength of visible light but instead is perceived when both blue and red wavelengths are present. Physiologically, magenta is the excitation of the long and short cone photoreceptors without a commensurate increase in activation of the medium cone photoreceptors. The linked-end spectrum is often called a color wheel (Figure 2.1b). Therefore, colors are a perceptual experience, while wavelengths of light are a physical property.

As you can see on the color wheel, there are smooth transitions between ad-

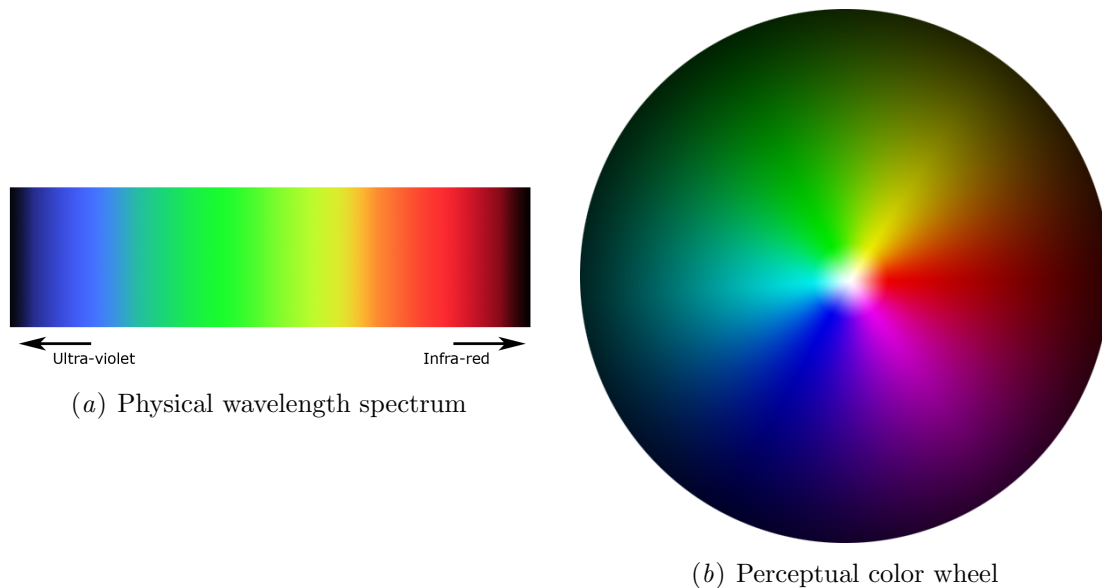


Figure 2.1: **Color perception: Wavelengths and wheels.** The visible light spectrum starts next to the infra-red region at red, then proceeds through all the colors until it smoothly transitions through blue on to violet and into the invisible ultra-violet region. Instead of this linear continuum, the human visual system creates a continuous wheel of color by connecting the red and blue colors with magenta.

adjacent colors, such as between green and yellow. But other colors are not adjacent, such as blue and yellow. It is not possible for blue to smoothly transition into yellow, without passing through a different color. These intermediate colors are made up of mixtures of different wavelengths of light, but how does the perceptual experience of a specific secondary color come about? Later chapters will show a model that makes significant progress at understanding this process, while the following sections explain what is known about the neural circuitry.

The following section discusses the major vision processing regions within biology, each simulated by the model.

2.2 The visual pathway

Much of the structure and function of the retina is well understood. As we move farther up the visual pathway, shown in Figure 2.2, the theories and data become less certain. Information about the Lateral Geniculate Nucleus (LGN) is much less definitive than the retina. And though much is known about the primary visual

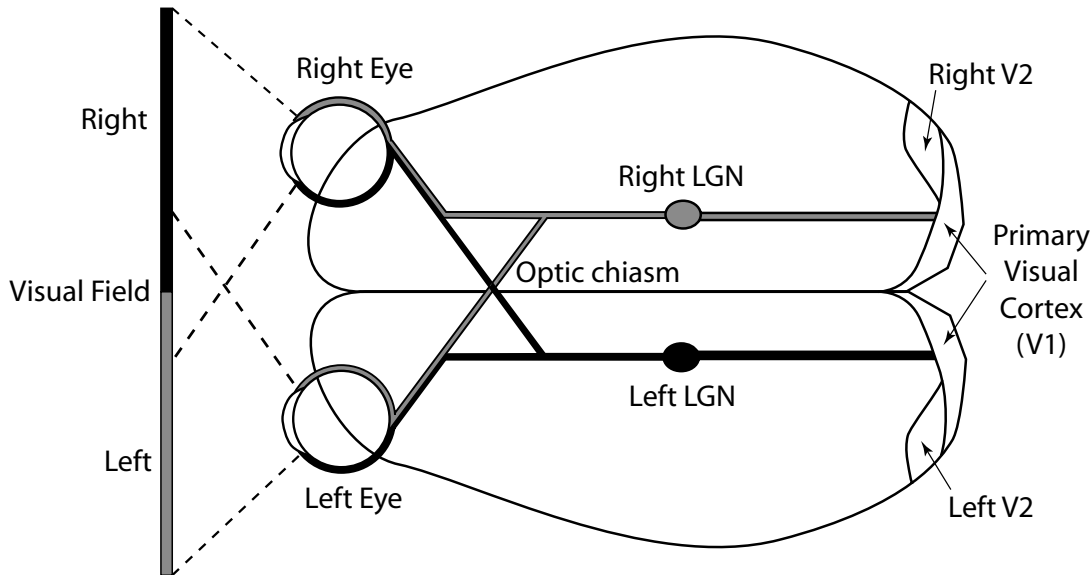


Figure 2.2: **The visual pathway.** The retinal response to light passes through the optic nerve into the lateral geniculate nucleus (LGN), and to the primary visual cortex. V1 then projects throughout the brain, including to V2. V2 is spatially close to, and connected with, V1. The large size of V1 (not visible in this top view) suggests a powerful role in vision. Figure adapted from Bednar (2002), with permission.

cortex (V1), only a few data points have been established for V2.

This section gives a short overview of the visual system starting at the retina and moving through to the higher visual areas, with each subsection emphasizing the physiology of color vision. Later chapters will describe how aspects of this color processing system have been implemented in the computational model.

2.2.1 Photoreceptors in the retina

Light enters the lens of the eye and is focused upon the retina, where the photoreceptors respond based upon the number of photons absorbed. There are two major classes of photoreceptors, called rods and cones. The cones are used for color and daytime vision, and rods are used for low-light (nighttime) vision. There are three types of cones separated into short, medium, and long based upon their relative light sensitivity functions (Brown and Wald, 1963). Long (L) wavelength sensitive cones are often called red-selective cones, medium (M) cones are called green-selective, and the short (S) are called blue-selective. The peak sensitivities of these cells do not map to what we would call pure red, green or blue, but the names are useful in describing




Figure 2.3: **Wavelength sensitivity functions.** RGB images are intended for display on a computer monitor. They can also be used as input activation for a simulated eye, but they need to be converted first, because the red and green activation for an eye has a much higher correlation than the red and green channels in a monitor. The top plot shows normalized long (red), medium (green), and short (blue) wavelength human cone sensitivity functions (Stockman and Sharpe, 2000; Stockman et al., 1999). The bottom plot shows RGB phosphor emission curves for a specific CRT monitor. The color of the curve specifies which phosphor was measured using a Sony GDM 500 (Wilson, 2004). These curves are typical for computer monitors.

their relative behavior (Itten, 1970). See the top plot in Figure 2.3 for the normalized sensitivity curves of L, M, and S cones as measured in humans.

Most computational vision simulations ignore color and assume an achromatic architecture. Creating training data is much simpler when working with achromatic simulations. Also, many lower animals do not have color vision, so doing a monochrome simulation is realistic. In trichromats such as macaque monkeys, achromatic neurons do exist, but a large fraction of cells are color selective.

The three types of cones described above are simulated by the model, as described in Chapter 3. Since this dissertation is modeling photoreceptors exposed to natural scenes, the model must have realistic retina activation patterns. Ideally, the raw light from nature would be used as input in the presented model, but it is not available. A mathematical transformation will instead be used to simulate photoreceptor activations using RGB images of natural scenes. The transformation simulates how the photoreceptors in a retina would respond if looking at a CRT (cathode ray

tube) monitor, and is presented in Appendix A (RGB to LMS transform). Alternatively, digital images from a calibrated camera can be used. With such calibrated images, the properties of the light-sensitive diodes in the camera are known precisely, and the images can be directly converted to photoreceptor activations.

As a shortcut, models can be trained on raw RGB images, but the results are not as biologically relevant as for models trained using the simulated photoreceptor activations. This is because an RGB picture of a natural scene has channel correlations that differ from what the biological eye sees. Figure 2.3 shows the difference between the light that CRT phosphors emit and the cone sensitivity functions of the human eye. The M and L cones overlap a great deal, yet there is little overlap in the G and R cones, and so models using RGB images without LMS transformation will not have realistic cross-channel correlations.

Simulating photoreceptor activations from RGB images can provide the majority of colors in the world, but note that some colors cannot be reproduced in this way. The color-space of a three-channel image is smaller than the color space of observable light, and therefore RGB images cannot produce every observable color (reviewed in Fairchild and Wyble (1998)). Even so, RGB images displayed on computer monitors are easy to use in the laboratory and for most experiments the shortcomings are negligible. The displays can be calibrated to show near lifelike images within the limits found in the camera optics, CRT phosphors, and ambient lighting of the testing environment. Since neuroscientists also use CRT monitors to measure brain activity, it is reasonable to expect that a model covering this color range will be sufficiently realistic for understanding biological data.

2.2.2 Retinal Ganglia and Lateral Geniculate Nucleus

The ganglion cells of the retina project to the Lateral Geniculate Nucleus, which in turn projects to the visual cortex, superior colliculus, and pretectum (reviewed in Gegenfurtner and Kiper, 2003; Kandel et al., 2000).

The response properties of retinal ganglion cells (RGC) are similar to those of the LGN. The spatially organized pattern of retinal photoreceptor cell activations that increases or inhibits the firing rate of a retinal ganglion cell or LGN cell is called its receptive field (RF). There are many studies that measure the receptive fields of cells in the retina and LGN (e.g. Conway, 2001; Gegenfurtner and Kiper, 2003; Lee, Kremers, and Yeh, 1998; Packer and Dacey, 2002). Receptive fields of RGC/LGN cells can be modeled using Gaussians with different variances and signs (Rodieck, 1965).

There are two main types of RGC/LGN receptive fields: center-surround, and coextensive. The receptive field for a center-surround LGN cell can be modeled by two Gaussians with opposite signs and different variances (Rodieck, 1965). Figure 2.4

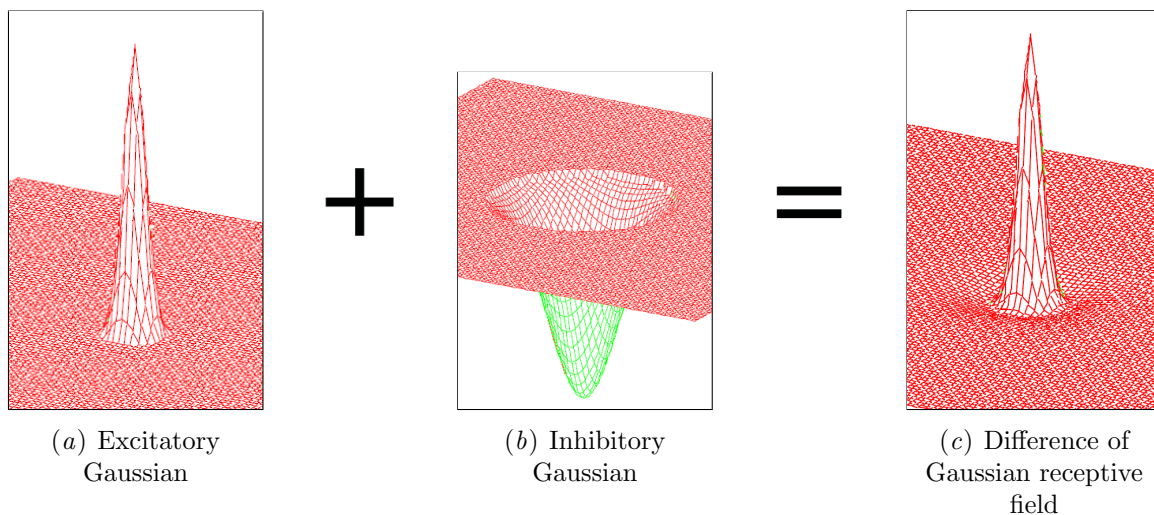


Figure 2.4: **Center-surround receptive field.** The flat surface in the 3D plot represents a zero effect on the activation of this neuron. High areas indicate areas where light will excite the neuron. Low areas are suppressive. The ON center-surround receptive field shown in (c) is modeled with the summation of the narrow positive Gaussian in (a) and the wide negative Gaussian in (b). A cell will activate strongly if bright light shines at the center of the cell's receptive field. If a light of the same intensity shines slightly offset from center, the cell will be suppressed.

shows an example of the center-surround receptive field of an RGC/LGN cell. The surface of each contour plot represents the contribution of a photoreceptor's activity to the excitation or inhibition of the LGN cell at that retinal location. In the figure, an increase or decrease in height from the baseline represents a proportional increase or decrease in activation of the RGC/LGN cell if there is illumination at that location. Figure 2.4a shows a very narrow (center) Gaussian at the center of the field, that will cause a very strong activation if light falls on the raised area. Meanwhile, Figure 2.4b shows a much broader negative (surround) Gaussian that will suppress the LGN cell if light is shown in the non-zero region. Figure 2.4c combines the excitatory and inhibitory Gaussians into a center-surround receptive field for an RGC/LGN cell. A bright light landing on the photoreceptors in the center of the cell's receptive field will activate it. If a light of the same intensity is shown slightly offset from center, then the cell will be suppressed instead. Since the center part of the receptive field in Figure 2.4c is excitatory, the cell is called an ON center-surround cell. There are also OFF cells that have the inverse of the receptive field in Figure 2.4; light in the center of the receptive field will suppress the cell, and light in the surround will excite it. Cells with these types of receptive fields can be found in the LGN parvocellular




Figure 2.5: **Red-green ON and OFF center-surround receptive fields.** Bright red areas in the figure denote the locations where red light will excite the cells. Darker red areas denote regions where red light will inhibit the cell. Similarly, green light will excite the cell in areas marked as light green, and inhibit the cells in dark green areas. Rows 1–4 show how neurons of these two types will respond to the given spots or rings of yellow light, which activates both red and green photoreceptors. The small yellow bars above each spike train show the onset and termination of the illumination, and the spike train shows a typical response. Figure adapted from (Kandel et al., 2000).

region (Gegenfurtner and Kiper, 2003), and in the retinal ganglion cells (Kandel et al., 2000).


Each cell can have different types of photoreceptors in its receptive field. LGN and retinal ganglion cells that receive stimulation from multiple types of photoreceptors are called luminosity sensitive cells, or wide or broad-band, and are achromatic. It is unclear whether luminosity cells only use long and medium photoreceptors in their receptive fields, or if they also include short-cone photoreceptors (Chatterjee and Callaway, 2002). Alternately, color opponent cells have connections to different types of cones for the excitatory and inhibitory areas of their receptive field. For example, an ON red-green color opponent cell will have an increase in activation if red light is shown in the center of the receptive field and/or green light is shown in the surround. Figure 2.5 gives further examples of how light will affect color opponent center-surround receptive fields.

The second major type of RGC/LGN receptive fields belong to spatially co-extensive (or spectrally opponent) cells. Spatially coextensive cells have a receptive field that lacks surround suppression. They are modeled with two Gaussians that have an equal variance, but with different signs and color preferences (Gegenfurtner and Kiper, 2003; Rowe, 2002). A cell with a blue-yellow coextensive RF will be excited when blue light is presented anywhere in the receptive field, and inhibited when yellow (red plus green) light is shown anywhere in the receptive field. Similarly, an OFF blue-yellow coextensive cell will be excited by yellow light shown in its receptive field and inhibited by blue light shown in its receptive field.


Red and green cones form center-surround receptive fields, while most research reports that blue-yellow selective cells have coextensive receptive fields (Calkins, 2001; Conway, 2001; Dacey, 2000; Rowe, 2002). However, Klug, Herr, Ngo, Sterling, and Schein (2003) suggest that a lens blurring effect may be causing blue selective center-surround cells to be misclassified as coextensive cells. There is also debate on whether OFF blue-yellow opponent cells exist (Klug et al., 2003), or if all blue-yellow cells are ON. Chatterjee and Callaway (2003) did find evidence for OFF blue-yellow cells but their role, if any, remains unclear. The model presented in this dissertation uses both ON and OFF blue-yellow coextensive cells, but these other possibilities have also been tested, and produced similar results.

2.2.3 Primary visual cortex


The neural response to the retinal stimulus passes through the LGN to the primary visual cortex (V1), located at the back of the brain (Figure 2.2). Optical imaging studies show that there are overlapping retinotopic maps in the V1 (Blasdel, 1992a,b; Blasdel and Campbell, 2001; Blasdel and Salama, 1986; Issa et al., 2001; Vanduffel et al., 2002). Each of these maps select for features of the visual stimuli such as



(a) Color selectivity map



(b) Orientation preference map



(c) Ocular dominance map

Figure 2.6: **V1 feature preference maps.** The V1 cortex contains overlapping feature maps for input properties like color, orientation, and eye of origin. All three plots were measured using optical imaging and show the same 5.3mm by 2.3mm section of macaque V1 (from Figure 8 in Landisman and Ts'o (2002b); reprinted with permission). Plot (a) is the color-selectivity map where color-selective regions are dark, and unselective regions are light. Plot (b) is the orientation preference map with the preferences of the each location colored according to the key to the plot's right. Plot (c) is the ocular dominance map with dark regions preferring the left eye, light regions preferring the right eye, and gray regions preferring both equally. The yellow dots in (c) show the location of the orientation pinwheels, and the color-selective regions are outlined in red. The pinwheels do not align with the color map, but color-selective regions tend to stay within the OD stripes.

orientation, eye of origin, spatial frequency, and color. This section first discusses some of the feature preference maps found in V1, and then reviews how these neurons connect to other V1 neurons. The discussion focuses on the aspects relevant to those presented in later chapters.

Figure 2.6 shows the color selectivity, orientation preference, and ocular dominance maps measured across the same 5.3mm by 2.3mm section of macaque monkey V1 cortex (Landisman and Ts'o, 2002b). Figure 2.6a shows a color-selectivity map where cortical regions that respond to colored retinal stimuli are dark, and light regions are unselective. Color-selective areas appear as blob-like patches in V1 and are found in the same locations that stain for cytochrome oxidase (Landisman and Ts'o, 2002b). Previous studies have suggested that color selective regions and cytochrome oxidase blobs were one and the same (Hubel and Livingstone, 1983; Hubener and Bolz, 1992; Sato et al., 1994). Landisman and Ts'o (2002a) showed that these cytochrome oxidase blobs do correlate with the imaged color patches, but that they

were not identical.

The same piece of V1 cortex can also be measured for orientation preferences, shown in Figure 2.6b. To measure orientation preference maps, the cortex is presented with oriented stimuli, and regions in the map are colored according to the orientation that causes the maximum activation. For example, in Figure 2.6b portions of the map colored yellow respond to vertical orientations, while blue areas prefer horizontal stimuli. Orientation maps have distinctive features called pinwheels, which are locations in the map where areas selective for all orientations come together at a single point.

A third type of feature map found in V1 (but not V2), is ocular dominance (or eye-preference) maps. Stimuli are alternately presented to the left and right eyes, to construct a map of how neurons in the cortex respond to stimuli in the left or right eye. Figure 2.6c shows an ocular dominance map where the map is black if the corresponding neurons activate to stimuli from one eye, white if the region responds to stimuli from the other eye, and gray if the cortex responds to stimuli in both eyes equally. Ocular dominance maps in macaque monkeys form stripes that smoothly transition between eye preferences.

On top of the ocular dominance map, Figure 2.6c has yellow dots that show the location of the orientation map pinwheels from Figure 2.6b, and also red outlines marking color-selective regions from Figure 2.6a. Landisman and Ts'o (2002b) did not find a relationship in registration or alignment between orientation column pinwheels and the color selective regions. However, they did find that color-selective blobs were often found in the monocular regions of ocular dominance stripes.

It may be that V1 is the first region where color selective cells begin to organize into the perceptual color space that is observed in psychophysical experiments. Xiao, Casti, Xiao, and Kaplan (2007) measured the color preferences of the color-selective regions and found that the cortex makes gradual transitions between preferred colors. For example, cortex that activates when green light is presented to the retina, is often located near cortex that responds to cyan (green plus blue) stimuli on the retina. Chapter 7 will present a model V1 that also has color preferences that make smooth transitions between perceptual colors.

Orientation selective neurons within the cortex form patchy lateral connections to other orientation selective neurons (Bosking et al., 1997; Sincich and Blasdel, 2001). Examples of these lateral connections are in Figure 2.7, which shows two orientation maps of tree shrew cortex. The small white plus signs in the plots are the locations of tracer injections, and the black dots are the axon terminals of the injected neurons. In Figure 2.7a the injected neurons are colored green and prefer near-vertical orientations. These neurons make nearby connections to areas with many different orientation preferences, but form long range connections to other vertical preferring neurons. In Figure 2.7B, the pattern is repeated: horizontal preferring neurons form




Figure 2.7: **Lateral connections of tree shrew V1 neurons.** V1 neurons form lateral connections to neurons with similar orientation preferences. Orientation maps of V1 cortex in the tree shrew are shown, with white plus-signs indicating the location of tracer injections, and black spots indicating axon terminals of the neurons. Plot (A) injected neurons that prefer orientations near vertical. These cyan colored neurons have connections to nearby neurons of all orientations, and have long-range connections to neurons with similar orientation preferences. Plot (B) shows how neurons that prefer orientations coded for red and purple (about 135°) again have close connections to neurons with a variety of preferences, and longer connections to neurons with similar orientation preferences. (Figure from Bosking et al. (1997). Copyright 1997 by the Society for Neuroscience, reprinted with permission.)

local connections to neurons of all orientations, but the longest connections are to other horizontal preferring neurons. These patterns are important because they play a role in modulating and controlling cortical responses, representing and associating information, and mediating development. It is not yet known how color-selective neurons are connected, but this connectivity might be crucial for color perception and the model in this thesis will make concrete predictions for lateral connections between color-selective neurons.

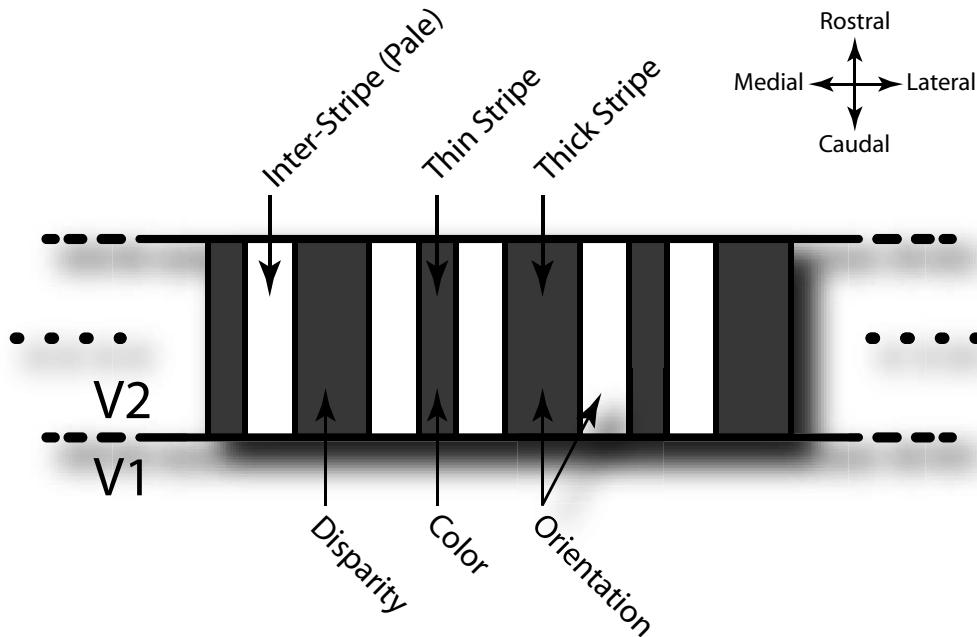


Figure 2.8: **V2 schematic organization.** The surface of V2 stains for cytochrome oxidase in a thin-pale-thick-pale stripe organization. Color selectivity is found in the thin stripes. Orientation selectivity is found in the pale and thick stripes, while disparity selective cells are located in the thick stripes.

2.2.4 Secondary visual cortex

The neurons in the primary visual cortex also form strong connections to the secondary visual cortex area (V2). Like V1, V2 neurons also form topographically organized connections and functional maps (Landisman and Ts'o, 2002b; Levitt, Kiper, and Movshon, 1994; Xiao, Wang, and Felleman, 2003). The surface of V2 can be divided into stripes that stain dark or light for cytochrome oxidase. The staining creates a repeating pattern of thin dark stripes, pale stripes (a.k.a. inter-stripes), and thick dark stripes (Roe and Ts'o, 1995; Shipp and Zeki, 2002a; Sincich and Horton, 2002a; Vanduffel et al., 2002; Yoshioka and Dow, 1996). These striped regions also show functional segregation, with color-selective cells appearing in the thin stripes, and orientation-selective cells located in the pale and thick stripes (Shipp and Zeki, 2002a). Unlike V1, ocular dominance maps do not appear in V2 (Ts'o, Frostig, Lieke, and Grinvald, 1990). These aspects of V2 physiology and functional organization are illustrated in Figure 2.8.

Xiao et al. (2003) used intrinsic optical imaging of V2 thin stripes to show that some color-selective regions respond to chromatic stimuli in the same spatial order as a standard perceptual color wheel. That is, a red-selective patch of cortex transitions into a yellow-selective patch of cortex, which transitions into a green-selective patch of cortex, and so on. The results presented in Chapter 7 also create perceptually organized bands of color and suggest how perceptual color space may be constructed from the mixing of color opponent LGN cells.

2.3 Developmental input-driven self-organization

The previous section reviewed the physiology of the human visual system. This section will present the developmental principles found in biology that form the foundation for the model presented in this dissertation. This section first discusses developmental studies supporting an input-driven self-organizing model of vision, and then presents the Hebb learning principle as a mechanism for the observed self-organization.

2.3.1 Visual cortex studies

During the 1960s Hubel and Wiesel broke new ground when they showed that maps in cat visual cortex changed based on what the animal saw during development (Hubel and Wiesel, 1962, 1963, 1967). If kittens had their eyes temporarily sutured closed at an early age, later sight was permanently impaired, or even completely blinded. If kittens were placed in environments that only contained vertical edges, then the orientation maps in the cortex were abnormally formed and the selectivity for orientations other than vertical were reduced (Blakemore and van Sluyters, 1975; Hirsch and Spinelli, 1970). These findings proved that properly formed adult feline cortex requires retinal stimuli with certain “natural” features during development. If abnormal stimuli are instead presented, then the self-organizing process will not form normal functional structures. Hubel, Wiesel, and LeVay (1977) went on to show that cortical self-organization in response to retinal stimuli also occurs in the higher mammals such as macaque monkeys. The self-organization of maps is a complex process, with some maps fully formed at birth, and some maps forming later. Color maps are likely to have a post-natal development phase (Sugita, 2004).

Abnormal cortical development is usually considered brain damage or unhealthy for the individual. Yet some abnormal variations may not be detrimental. Jordan and Mollon (1993) report that they found women who are tetrachromats (i.e. four types of cones), with at least one woman having superior color perception. This findings suggests that the visual cortex is adaptable and can make use of additional sensory information. Similarly, Jacobs, Williams, Cahill, and Nathans (2007) has

shown that mouse cortex can also use increased photoreceptor selectivity to achieve superior color perception. They genetically modified a mouse by replacing one photopigment of mouse photoreceptors with the DNA of a human photopigment, thereby increasing the color discrimination of the mouse. Overall, the evidence available suggests that cortical color processing circuitry may develop automatically from features of the visual environment, as processed by subcortical circuitry.

2.3.2 Synaptic plasticity and the Hebb principle

In 1949, Donald Hebb proposed a learning principle between neurons that also works as a mechanism for the self-organization of neural networks (Hebb, 1949). The adaptation mechanism, later called Hebbian learning, has become a fundamental model within neuroscience and has been used to explain long term potentiation (LTP) as well as everything from memory (Horn, Levy, and Ruppin, 1998; Tsien, 2000) to vision (Lee, Eglon, and Wong, 2002a) to consciousness (Miranker, 1997).

Hebbian learning is based upon the premise that connections between neurons should be strengthened when they are active simultaneously. When a neuron is activated, the learning will cause other neurons which are also active to play a greater role in exciting the neuron in the future. One form of the Hebbian learning rule (adapted from Haykin, 1999) is,

$$\Delta\omega_{ij} = \alpha c_i d_j, \tag{2.1}$$

where $\Delta\omega_{ij}$ is the amount of the weight adjustment for the connection weight ω_{ij} which connects neuron c_i to the input d_j , and the parameter α is the learning rate which can be changed depending upon the need.

2.3.3 Homeostatic plasticity

If there was no balancing mechanism to temper long term potentiation then connections and firing rates would grow without bound. A number of biological mechanisms help to regulate long-term and short-term potentiation, to keep neurons firing within useful ranges (Bourgeois, Jastreboff, and Rakic, 1989; Pallas and Finlay, 1991; Turrigiano, 1999). An effective model of these regulating mechanisms is to use weight normalization. That is, to keep the total value of all weights to a neuron fixed, and to change the connection weights relative to each other, so that a connection weight can be strengthened but at the cost of another weight being proportionally weakened. Hebbian learning with normalization is a model that encapsulates both the LTP and regulating mechanisms within the brain, and will be used later for the learning mechanism within the vision model.

With this biological background in place, the next section presents computational models of self-organizing maps based upon the visual cortex.

2.4 Self-organizing models of the visual cortex

The first self-organizing computer model of the visual cortex was run in 1973, when von der Malsburg simulated V1 neurons on a UNIVAC running at 1 MHz (von der Malsburg, 1973). His 2D cortex represented each cortical column as a single unit, and included fixed excitatory and inhibitory lateral connections. Inputs were oriented bars of light represented as matrices of zeros and ones, and weights were adjusted using Hebbian learning. The computer program created orientation-selective neurons reminiscent of mammalian cortex, and revealed the promise that the fledgling computer technology offered to brain researchers.

2.4.1 Monochrome vision models

The next step in the development of self-organizing cortex models was the development of self-organizing feature maps (SOMs). These maps are mathematical systems that arrange high-dimensional information into organized geometric patterns following regularities contained within the input data and the structure of the network (Kohonen, 1982a,b, 2001). SOMs develop features that have similarities to biological cortex (Durbin and Mitchison, 1990; Goodhill, 1993; Obermayer, Ritter, and Schulten, 1990a,b; Piepenbrock, Ritter, and Obermayer, 1997). Orientation maps tend to be the dominant map studied, followed by ocular dominance (Miller, Erwin, and Kayser, 1999).

These models save computer time and memory by not explicitly including individually adjustable lateral connections, which also makes the resulting models easier to analyze mathematically. However, explicit lateral connections increase biological realism, and extend the predictive power of the computational model because the model's connections can be compared to the neuronal connections in biology (Bartsch and van Hemmen, 2001; Kalarickal and Marshall, 1998; Sirosh and Miikkulainen, 1994).

Alternatives to SOMs include correlation-based learning models (CBLs). If the visual system is assumed to be linear, then simple functions can be used to represent the long-term correlations of the inputs (Miller, 1994). CBL models create reasonable orientation maps, but fail to form some biologically common features such as linear zones, which are found in competing models using Hebbian learning (Miikkulainen et al., 2005; Obermayer, Sejnowski, and Blasdel, 1995).

The most complete self-organizing cortex models blend SOMs, lateral cortical

connections, and non-linear Hebbian learning (Alexander, Bourke, Sheridan, Konstantatos, and Wright, 2004; Barrow and Bray, 1993; Burger and Lang, 1999, 2001). The earliest system to include these components and show patchy lateral connections was LISSOM (the Laterally Interconnected Synergetically Self-Organizing Map; Miikkulainen et al., 2005; Sirosh, 1995; Sirosh and Miikkulainen, 1994). This model showed that patchy lateral connections help reduce redundancy within the cortex, and may also aid in feature detection.

The LISSOM model has proven powerful, robust, and flexible, and it has been extended to simulate properties other than the more traditional orientation and ocular dominance maps. These additional features include motion, disparity, spatial frequency, perceptual grouping, and even face-selective cortical areas (Bednar, 2002; Bednar and Miikkulainen, 2003; Choe, 2001; Palmer and Bednar, 2006; Ramtohl, 2006). The model even explains the tilt aftereffect and the McCollough effect optical illusion, providing further evidence that the self-organization theory accurately describes part of the biological vision system (Bednar and Miikkulainen, 2000; Ciroux, 2005).

2.4.2 Models of color selectivity

When modeling the visual pathway from the retina up to V1, few self-organizing computational models include color stimuli. There are at least four main reasons: First, there is much to learn about vision without needing to include color. Second, many of the biological imaging studies being modeled are performed on animals that do not have full trichromatic color vision. Third, memory and computer processing power have been insufficient for the added complexity that color requires. Fourth, color vision is complex to model and requires good software support.

Several single neuron models of color and color constancy have been published (Courtney, Finkel, and Buchsbaum, 1995; Stanikunas, Vaitkevicius, and Kulikowski, 2004; Usui, Nakauchi, and Miyake, 1994; Usui, Nakauchi, and Miyamoto, 1992). However, these models do not work on the map level and therefore cannot be used to study how color maps form and interact with other feature maps.

Saarinen and Kohonen (1985) applied SOMs to create color maps in a modeled cortex years before SOMs were being routinely used to model orientation maps. These maps formed bands of selectivity that made gradual transitions but did not resemble the biological cortex. However these maps are still remarkable because they predated biologically plausible color maps by a decade.

Barrow, Bray, and Budd (1996) presented a model with an architecture similar to the one in this dissertation, with photoreceptor, lateral geniculate nucleus, and V1 sheets. Their simulation develops color blobs and is consistent with the results in this paper, but it is dichromatic (with no blue photoreceptors), has a single eye, and has

fixed, isotropic lateral connections. Thus it cannot address the issues studied in this thesis, such as the interactions between ocular dominance (OD), orientation (OR), and color maps, and how the lateral connections interact with the structure of each map.

A V1 model by Nakagama and Tanaka (2004) develops patterns similar to cytochrome oxidase blobs. Their cortical neurons are not orientation selective, but they have included two retinas to create ocular dominance stripes. They find that the cytochrome oxidase blobs remain within the ocular dominance bands, as they do in biology (Bartfeld and Grinvald, 1992; Landisman and Ts'o, 2002b). They are able to manipulate the position of cytochrome oxidase blobs by adjusting between-eye correlation, but they use artificial stimuli, and do not model all three types of photoreceptors.

Rao, Cecchi, Peck, and Kozloski (2005) have designed a model for color categorization. Similar to the model presented in this paper, their model is a trichromatic simulation that creates a self-organized color map. However, they do not model either the retina and LGN regions or cortical lateral connections, and do not consider how the color map interacts with the orientation and ocular dominance maps. Thus, their model is a more abstract simulation of color processing in general, rather than a specific model for how color-selective neurons are organized in V1.

The LISSOM simulation presented in this paper goes beyond the Nakagama and Tanaka model or the Barrow and Bray model to include self-organized orientation and ocular dominance selectivity with trichromatic inputs to undifferentiated cortical neurons. This simulation is the first to bring together each of these different input features into a single computational model. Such a comprehensive model leads to unique interaction results that are not possible with simpler simulations. In particular, the new model makes it clear why color is organized differently from orientation and ocular dominance.

2.5 Conclusion

This chapter reviewed the biology and functionality of the early human visual system. There is a feed-forward information pathway starting at the retinas and moving up to the secondary visual cortex. Research suggests that the cortex self-organizes topographic maps based upon the input stimuli received from the retina, and that Hebbian learning is a useful model for this self-organizing process.

The next chapter will present the details of the trichromatic LISSOM model that is used to generate the results presented in this dissertation.

Chapter 3

The LISSOM model

This dissertation extends LISSOM to include trichromatic photoreceptors in the retina and color-opponent cells in the LGN, which then allows color-selective regions to form in V1. Adding color is a significant extension to the LISSOM model since the color inputs add many new dimensions of complexity. In this section the architecture of the model is presented first, with detailed descriptions of each neural region. Second, the activation, settling, and learning equations are discussed. Third, the network training stimuli are described. Thus, the chapter gives a complete description of the computational platform used in later chapters.

3.1 Architecture

The model (shown in Figure 3.1) consists of 24 two-dimensional sheets of artificial neurons representing different topographically organized sets of neurons: six sheets of retina photoreceptors representing two eyes, each with three different cone types; 16 paired sheets of ON-center/OFF-surround or OFF-center/ON-surround LGN units, each with a different color combination; one sheet of neurons representing V1; and one sheet of neurons representing V2. Images of natural scenes are presented to the simulated retinas which then activate the LGN regions, which activate V1, and in turn activate V2.

The input to the simulation consists of a series of color bitmap images separated into long, medium, and short photoreceptor activation values. The LGN units represent the entire processing pathway between the photoreceptors and V1, including the retinal ganglion cells. Four common types of ganglion cell receptive fields are modeled: luminosity, long-center/medium-surround, medium-center/long-surround, and short/(long and medium). Because the focus is on the two-dimensional organization of the cortex, each cortical neuron corresponds to a vertical column of cells

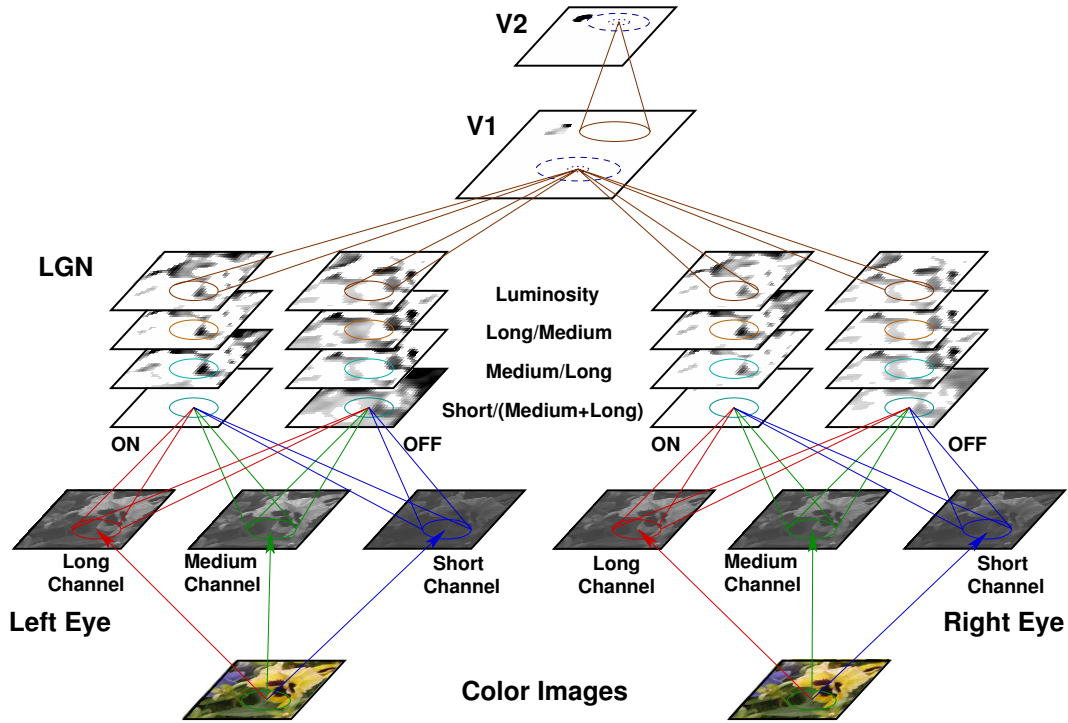


Figure 3.1: **LISSOM model of orientation, long/medium/short color selectivity, and ocular dominance.** The model is a hierarchy of sheets of neural units, modeling the visual pathway from the retinal photoreceptors to V2. In this figure, the model is being presented with a close-up image (by Olmos and Kingdom, 2004) of a yellow flower with brown spots and the edge of a blue flower. For training and activation, calibrated natural images are converted into long-, medium- and short-cone photoreceptor activation values. These values become the bitmap activation patterns of the photoreceptor cells in the model. The units in the LGN sheets have Difference of Gaussian and coextensive receptive fields projecting from the retina layer. V1 and V2 neurons also have lateral excitatory (small dotted circle) and lateral inhibitory (large dashed circle) connections to their neighbors. Through training, neurons in the V1 and V2 sheets develop specific feature preferences, forming into maps for color, eye and orientation preference.

through the six anatomical layers of the cortex.

3.1.1 Eyes

The photoreceptors described in Section 2.2.1 are modeled with three 2D matrices of floating point numbers. Each matrix represents one type of photoreceptor, and each floating-point number in the matrix represents the activation (i.e., graded voltage) of a single photoreceptor. The retina can “see” images if the matrices are filled with the simulated activation level of photoreceptors looking at a scene. As reviewed in Section 2.2.1, RGB images do not have the same inter-channel correlations as photoreceptor activations but they can be used as training images if they are first converted. Section 3.3 and Appendix A discuss how to calculate photoreceptor activation by converting red, green, blue (RGB) image measurements into simulated long, medium, short (LMS) activations.

3.1.2 Lateral geniculate nucleus

Each model “LGN” unit represents all of the computational processing between the photoreceptors and V1, including the retinal ganglion cells. Each LGN unit receives input from photoreceptor units within circular receptive fields (RFs). Figure 3.1 shows the connections for four example LGN units. All units at the same position in the other 12 LGN sheets in the figure receive input from the same set of photoreceptors but with different connection weights. Electrophysiological studies have shown that the pathway to V1 contains ganglion cells with these types of receptive fields (Chatterjee and Callaway, 2003; Dacey, 1994; Livingstone and Hubel, 1984).

Similar to the biological naming conventions described in Section 2.2.2, if light falling on the center of a receptive field activates a cell, then it is labeled an ON LGN receptive field. Likewise, cells with OFF receptive fields are inhibited when light falls on the center. LGN cells that have long and medium cone inputs are modeled with center-surround receptive fields. Center-surround receptive fields use difference of Gaussian (DoG) connections to the photoreceptors, as previously described in Section 2.2.2. A DoG is composed of two Gaussian activation functions of opposite sign, with one excitatory Gaussian and one inhibitory Gaussian. The two Gaussians have different variances but identical centerpoints; the smaller Gaussian determines the properties of the center, and the larger one determines those of the surround (Lee et al., 1998; Packer and Dacey, 2002). In contrast, modeled LGN cells with short-cone inputs have coextensive receptive fields, i.e. receptive fields that use center-sized Gaussians for both the excitatory and inhibitory fields.

Each LGN sheet is sensitive to a different set of photoreceptors. Some LGN cells combine input from all three types of photoreceptors to create broad-band cells.

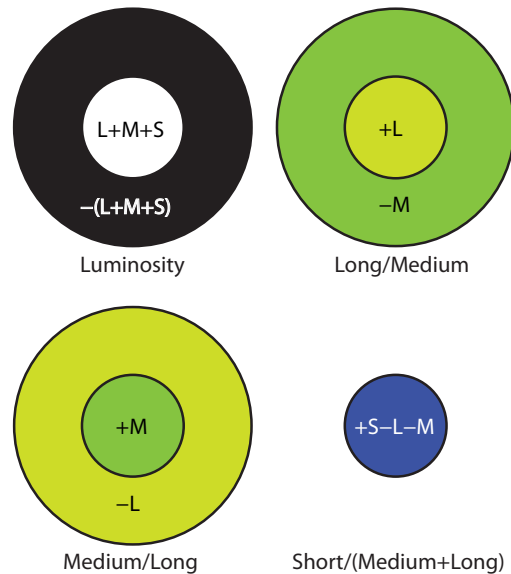


Figure 3.2: **Sample receptive fields for units in LGN sheets in the LISSOM model.** The spatial organization of four modeled receptive fields are shown in this plot. The model includes four more RF types identical to the above except for sign: $-(L+M+S)/(L+M+S)$, $-L/+M$, $-M/+L$, and $-S+L+M$. Using the Long, Medium, and Short photoreceptor sensitivity functions means that LGN receptive fields are maximally selective for colors other than the pure perceptual colors. The center and surround receptive fields in the figure have been filled with the approximate color that will cause the peak response of the indicated sign, in that part of the receptive field (RF). Cells with these types of receptive fields are regularly found in macaque.

Other cells have connections to different types of cones for the ON and OFF regions of its receptive field, thus forming color-opponent cells. For example, a red/green color-opponent cell can be created by combining activation from an ON-center receptive field that is selective to long wavelengths, and an OFF-surround receptive field that is selective to medium wavelengths. Figure 3.2 shows four of the eight different receptive field types found in the model LGN. The four not shown are identical to the four RFs are shown except that they have opposite signs. These same cell types have been found experimentally in the parvocellular layers of the LGN (Gegenfurtner and Kiper, 2003), and in the ganglion cell layer of the retina (Kandel et al., 2000).

3.1.3 Primary visual cortex

Just as LGN neurons receive input from the retina sheets, V1 neurons receive inputs from circular connection fields on all 16 LGN sheets (Figure 3.1). Each V1 neuron has connections to each LGN sheet, with the connection strengths chosen randomly. Each V1 neuron also has lateral excitatory and inhibitory connections in a circular radius around it as shown in Figure 3.3, with the weights initially having a Gaussian distribution. The connection weights are eventually learned in a Hebbian adaptation process as described in Section 3.2.

A single V1 neuron represents a vertical column of cells, which is a viable approach because biological response properties tend to be similar in each vertical penetration. The resulting model is thus a 2D sheet where the layered structure has been collapsed into a single unit representing each column.

3.1.4 Secondary visual cortex

The secondary visual cortex is modeled with another 2D sheet of LISSOM neurons. Each V2 neuron has a circular receptive field projecting from the V1 sheet. The radius of the receptive field gives each V2 neuron double the retinal visual field of V1 neurons. Like V1, each V2 neuron also has lateral excitatory and inhibitory connections in a circular radius around it as shown in Figure 3.3. All connections are eventually learned through a Hebbian learning process.

3.2 Running the model

In LISSOM, lateral excitatory connections have a short range and inhibitory connections connect to a much wider region. In biological cortex, the longest connections in V1 and V2 are excitatory, but their net effect is inhibitory at high contrasts due to strong local inhibitory interactions with a high threshold for activation (Grinvald, Lieke, Frostig, and Hildesheim, 1994; Hata, Tsumoto, Sato, Hagihara, and Tamura, 1993; Hirsch and Gilbert, 1991; Weliky, Kandler, Fitzpatrick, and Katz, 1995). The model uses explicit long-range inhibitory connections for simplicity, which is a valid approximation because the high-contrast portions of each image have the largest influence on Hebbian learning (Bednar and Miikkulainen, 2000). LISSOM-like models that use long-range excitatory connections with explicit inhibitory interneurons behave similarly (Law and Bednar, 2006).

For each training input image, the model units are activated and adapted in a three-stage process:




Figure 3.3: **LISSOM model.** The LISSOM model is an artificial neural network with a 2D sheet of neurons. Each neuron has three types of connections: afferent input projections coming from a spatially organized input surface, lateral short-range excitatory connections, and longer-range lateral inhibitory connections. All of the connection weights are learned through Hebbian self-organization from an initial random starting point. For the V1 map, the Network sheet is V1 and the Input Surfaces are the LGN sheets; for the V2 map, the network sheet is V2 and the Input Surface is V1.

1. **Feedforward Activation:** Find the initial activation by summing all the incoming connections.
2. **Lateral Settling:** Iteratively settle the initial activation using the lateral connections.
3. **Learning:** Adjust the afferent and lateral connection weights through normalized Hebbian learning.

The next three subsections present the equations governing each stage.

3.2.1 Activation

For each input image, the activity level for each LGN, V1, and V2 cell is calculated. Each LGN unit (i, j) computes its input activation s_{ij} as a scalar product of a fixed weight vector and an input vector for each of its connection fields:

$$s_{ij} = \sum_{\rho \in C} \gamma_{\rho} \sum_{\rho ab} \xi_{\rho ab} A_{\rho ab, ij}, \quad (3.1)$$

where ρ is an index over the set C of connection fields (e.g. $C = \{\text{long-center, medium-surround}\}$ for an on-center long/medium cell) of neuron (i, j) , $\xi_{\rho ab}$ is the activation of neuron (a, b) in connection field ρ , $A_{\rho ab, ij}$ is the weight value for the connection to $\xi_{\rho ab}$, and γ_{ρ} is a constant scaling factor for each connection field. The LGN cell activity is then $\sigma(s_{ij})$, where σ is a piecewise linear sigmoid activation function, with zero output below a lower threshold δ , 1.0 output above an upper threshold β , and a linear increase between the thresholds. Both δ and β are set at the start of training and then adjusted according the schedule in Table B.4.

3.2.2 Cortical settling

Each V1 and V2 neuron computes its initial input activation like an LGN cell using Equation 3.1, but then the neuron activation settles through short-range excitatory and long-range inhibitory lateral interaction:

$$\eta_{ij}(t) = \sigma \left(s_{ij} + \gamma_E \sum_{kl} \eta_{kl}(t-1) E_{kl, ij} - \gamma_I \sum_{kl} \eta_{kl}(t-1) I_{kl, ij} \right), \quad (3.2)$$

where $\eta_{kl}(t-1)$ is the activity of another cortical neuron (k, l) during the previous time step, $E_{kl, ij}$ is the excitatory lateral connection weight on the connection from that neuron to neuron (i, j) , and $I_{kl, ij}$ is the inhibitory connection weight. All connection weights have positive values. The scaling factors γ_E and γ_I are the relative strengths of excitatory and inhibitory lateral interactions, which determine how easily the neuron reaches full activation.

The V1 activity pattern starts out diffuse, but within a few iterations of Equation 3.2, it converges into a small number of stable focused patches of activity bubbles. The V1 settled activity pattern is the input used for V2.

3.2.3 Learning

After activity in the cortex model has settled, the connection weights of all V1 and V2 neurons are modified. All weights adapt according to the Hebb rule, normalized so that the sum of the weights from each sheet is constant for each neuron (i, j) :

$$w'_{pq,ij} = \frac{w_{pq,ij} + \alpha X_{pq} \eta_{ij}}{\sum_{uv} (w_{uv,ij} + \alpha X_{uv} \eta_{ij})}, \quad (3.3)$$

where η_{ij} stands for the activity of neuron (i, j) in the final activity bubble, $w'_{pq,ij}$ is the new connection weight and $w_{pq,ij}$ the current connection weight, α is the learning rate for each type of connection (α_A for afferent weights, α_E for excitatory, and α_I for inhibitory), and X_{pq} is the presynaptic activity. The larger the product of the pre- and post-synaptic activity $X_{pq} \eta_{ij}$, the larger the weight change. Before training, cortical neurons are connected to all nearby neurons. But at long distances, few neurons have correlated activity and most long-range connections eventually become weak. The weakest connections are eliminated periodically, resulting in patchy lateral connectivity similar to that observed in biological visual cortex.

Initially, the learning rate of the connection weights and activity blobs are large and make strong changes to the cortex. But as the simulation progresses, the learning rate decreases so that fine details appear in the cortex and activity blobs become smaller. The learning rate is controlled by a small number of free parameters in the model according to a fixed parameter modification schedule. Appendix B contains the parameter modification schedules for the results presented in this dissertation which were empirically found through repeated experimentation.

3.3 Training stimuli

The equations from the previous section allow the connection weights to be self-organized, given a set of input patterns. The type and order of the training stimuli strongly affects the final organization of the neural network. In order to create realistic visual cortex maps, the simulation needs realistic retinal activation patterns. Therefore, color digital images of natural scenes were used as training input.

The training images were selected from freely available databases of foliage and landscape images, explicitly eliminating photographs that contained man-made objects that might have unnatural color statistics. The specific results presented in this dissertation were based on a subset of the McGill images (Olmos and Kingdom, 2004). These images come from cameras that directly convert the responses of the light-sensitive diodes in the camera into the long, medium, and short (LMS) photoreceptor activation values of the retina, and are called calibrated images.

Training the model using calibrated images is slightly more accurate than using the RGB-to-LMS transformation reviewed in Chapter 2. The difference arises because the camera-to-photoreceptor mapping is known for a calibrated image but the RGB-to-LMS conversion can only estimate the original natural light. In addition, any rounding errors in the original RGB image are further compounded. However, the difference in accuracy is not crucial for LISSOM. The McGill corpus is used in this thesis, but as described in Section 7.1, other image datasets using the RGB-to-LMS transformation give similar results.

V1 regions are trained first, using 20,000 random patches of 110×110 pixels sampled from the input images. Then V2 regions are trained using an additional 20,000 random patches. For each iteration of the simulation a random patch of a random image is selected and presented to both eyes. In order to provide eye differences that can drive ocular dominance, during V1 training the brightness of the patch in the left eye is chosen randomly in the range 0 to 1.0, and the brightness of the corresponding patch in the right eye is then chosen so that the total brightness is constant. During V2 training the brightness of the left and right eyes are equal. Other types of differences between the two eyes (such as position) can also be used to drive ocular dominance development (Miikkulainen et al., 2005), but these have not yet been investigated in the present model.

3.4 Conclusion

This chapter presented a trichromatic computational model of primate vision. The model has a hierarchical organization and uses lateral connections and Hebbian learning to self-organize cortical regions. To get realistic biological cortical maps, the training input to the model consists of images of natural scenes. To help understand how natural images affects the self-organization of the model, the next chapter presents several statistical analyses of the training images, as well as some predictions of how the images will affect the developed cortex.

Chapter 4

Statistics of natural images

The previous sections have described a highly complex self-organizing biological system that processes what is seen by the eye. So what do eyes see in the world? They see both natural and man-made objects. Human artifacts have many more right edges and colors than do the natural scenes from a forest or a landscape. This dissertation focuses on natural objects since the eye evolved while seeing only natural scenes. There is a large body of literature that studies natural images from a psychophysical or psychological perspective. Unfortunately, most of the work does not study the relationships among photoreceptor types which is most relevant to this research. To fill that void, five numerical analyses were performed on image corpora to study the color information within photoreceptor color channels.

If a simulation is trained on images taken of the natural world, and these images have been characterized in some quantitative manner, then conclusions can be made about how the statistics of the natural world affect the self-organization of the model. Numerical analyses of image corpora are presented here in preparation for studying how natural images affect the self-organization of a modeled cortex. These studies resulted in five observations. First, activations of the three photoreceptor types are extremely correlated, making unique channel information difficult to extract. Second, simulated photoreceptor activations have between two and five bits of entropy per pixel. Third, individual images tend to have few unique colors. Fourth, the principal component of small image patches is typically an orientation-selective receptive field. Fifth, natural images tend to have more luminosity transitions than color transitions. Each of these observations, and how they affect the model, will be discussed in the following sections.

4.1 Correlations between photoreceptor types

As discussed in Section 3.3, the training corpora for this dissertation come from RGB images of natural scenes that have been converted to long, medium and short (LMS) photoreceptor activations either by using a calibrated camera or by applying the transformation in Appendix A. Within this section, a channel refers to all the pixel values of the red, green, or blue pixels of an RGB image, or the activation values of the long, medium, or short photoreceptors in the LMS transformation of an RGB image.

Simulated retinal activations are highly correlated because the long, medium, and short photoreceptor sensitivity functions overlap. Pearson’s correlation coefficient was calculated to measure these correlations. The long, medium, and short photoreceptors would have a correlation of 1.0 if they all had the same value for each wavelength, and have a lower correlation if they are different. Table 4.1 shows the source and sizes of different image sets obtained from a number of geographic locations. Though there are sets of images from Texas (De Paula, 2007a,b; Visual Delights, 1999), Canada (Olmos and Kingdom, 2004) and even Africa (Lovell, Tolhurst, Párraga, Baddeley, Leonards, and Troscianko, 2005; Troscianko, Párraga, Leonards, Baddeley, Troscianko, and Tolhurst, 2003) all images have similar correlation coefficient values. The different corpora mostly contain images of mountains, trees, bushes, fruit, water, sky, flowers, snow, and rocks. The minor variations in correlation come from the content of the source images. For example, the Winter image set contains snow-covered scenery that is mostly white, whereas the Flowers corpus contains zoomed photographs of vibrantly colored flowers. As a result, the lowest r^2 correlation for Winter is 0.995, whereas the Flowers corpus is less correlated with a minimum of 0.784. The r^2 correlation values for all the image corpora are shown in Figure 4.1.

It is remarkable that with the wide variety of sources and the selection of objects within the images, that the channel correlations are consistently so high. In each case, the long and the medium photoreceptor activations have the highest correlation. This relationship is not surprising given that the long and medium photoreceptor sensitivity curves have such a high overlap. Chapter 8 discusses in detail how channel correlations affect the modeled cortex.

The correlation coefficient gives a good single number to compare training sets, but the value does not illuminate how much actual information is available in the stimuli, and how much redundancy is found in the channels. This topic will be presented in the next section.

Name	Arboretum	Austin Vicinity	Flowers	Foliage	Fruit Kibale, Uganda	Landscapes	Winter	
Source	Personal	Visual Delights	Personal	McGill	McGill	Troscianko	McGill	
Images	31	74	32	136	16	30	34	
L:M r^2	0.997	0.997	0.978	0.987	0.978	0.981	0.992	0.999
Std Dev	0.003	0.007	0.031	0.017	0.027	0.027	0.010	0.000
Mean Error	0.001	0.001	0.005	0.001	0.007	0.005	0.002	0.000
L:S r^2	0.921	0.895	0.784	0.906	0.877	0.855	0.901	0.995
Std Dev	0.062	0.125	0.165	0.078	0.098	0.076	0.105	0.001
Mean Error	0.011	0.015	0.029	0.007	0.025	0.014	0.018	0.000
M:S r^2	0.939	0.910	0.796	0.931	0.910	0.897	0.937	0.996
Std Dev	0.047	0.098	0.148	0.064	0.074	0.060	0.070	0.001
Mean Error	0.008	0.011	0.026	0.005	0.019	0.011	0.012	0.000

Table 4.1: **Correlations between channels in various image corpora.** This table lists the names and sources of the primary image sets used in the dissertation. Each image set was converted into simulated photoreceptor activations, and then the pixel-level correlations between the three types of photoreceptors (long, medium, and short) were measured. Each column in the table shows the number of images in the corpus, the average pair-wise channel correlations of those images, the standard deviation, and the standard mean error. The correlation between the long and medium channels is generally the highest in the image sets because these wavelength sensitivities overlap significantly.

4.2 Independent and shared information between photoreceptor types

It has been established that the individual photoreceptor channels are highly correlated. By using information theory it is possible to calculate exactly how much information is unique to each long, medium and short channel, and how much information is shared. Measuring the information in a training input will establish what information content thresholds are necessary for the model to self-organize properly.

4.2.1 Joint entropy

The total entropy within two channels is calculated using the joint entropy equation (Yao, 1993)

$$H(X, Y) = - \sum_{x,y} p_{x,y} \log(p_{x,y}), \quad (4.1)$$

where x and y are the values of the channels X and Y during a single event. The probability of the two channels taking on the specific values x and y simultaneously is




Figure 4.1: **Graphical comparison of channel correlations in Table 4.1** The Winter corpus has the highest correlations because the images are dominated by white snow. The Flowers corpus has the lowest long:short and medium:short correlation values, and ties Fruit for the lowest long:medium value. Bright colors in the Flowers corpus causes the channels to have lower correlations than the other corpora. The error bars show the standard mean error of the images within each corpus.

denoted as $p_{x,y}$. The joint entropy for each pair of channels is measured by assigning X and Y first to the long and medium channels, then the long and short channels, and finally the medium and short channels.

The joint entropy equation is an extended form of the more common single-channel entropy equation which will be used later for the minimum unique entropy calculation and the redundancy calculation. To calculate the entropy of a single channel, eliminate Y from Equation 4.1, and use p_x instead of $p_{x,y}$. To calculate the joint entropy of all three channels, add a third channel Z and use $p_{x,y,z}$.

In the image corpora, the joint entropy of the L, M, and S channels tend to be between 13 and 16 bits of information per LMS pixel, as shown in Table 4.2, with a theoretical maximum of 24 bits. The joint entropy of any two channels is between 9 and 13 bits of information per pixel. The Fruit corpus has the most entropy of all

Corpus Name	L:M	L:S	M:S	L:M:S
Austin Arboretum	10.12	11.95	11.88	13.98
Austin Vicinity	9.77	11.72	11.74	13.81
Flowers	10.39	11.59	11.37	14.17
Foliage	11.22	12.48	12.23	15.75
Fruit	12.07	12.97	12.70	16.30
Kibale, Uganda	9.75	11.22	11.06	13.59
Landscapes	11.38	13.96	12.83	15.57
Winter	10.78	12.09	12.03	14.03

Table 4.2: **Joint entropy of channels in various image corpora.** Joint entropy measures how much information is contained within pairs or triples of information channels. The long:medium pair consistently has less information than the other channel pairs because the long and medium sensitivity functions contain redundant information. The Fruit corpus has the largest per-pixel entropy, which means it contains the most information of all image sets.

Corpus Name	L:M	L:S	M:S	L:M:S
Austin Arboretum	4.35	2.22	2.30	1.44
Austin Vicinity	4.40	2.10	2.09	1.50
Flowers	1.27	1.03	0.66	-0.04
Foliage	3.50	1.96	2.11	1.57
Fruit	3.05	1.97	2.05	0.97
Kibale, Uganda	3.78	1.69	1.81	1.22
Landscapes	3.66	2.03	2.20	0.88
Winter	4.57	3.30	3.26	2.14

Table 4.3: **Mutual information of channels in various image corpora.** The corpora contain widely varying levels of shared information. Austin Vicinity and Winter have the highest mutual information between the long and medium channels. Meanwhile the Flowers corpus has the lowest shared information among its channels. Low mutual information and a high entropy may help a neural network to efficiently self-organize channel-specific selectivity.

the corpora with 16.30 bits per LMS pixel. The joint entropy reveals that going from two to three channels of information provides a mere two or three bits of additional information.

Corpus Name	L:M	L:S	M:S
Austin Arboretum	2.88	4.72	4.64
Austin Vicinity	2.68	4.63	4.65
Flowers	6.43	6.40	6.77
Foliage	3.81	5.07	4.92
Fruit	4.42	5.32	5.23
Kibale, Uganda	2.96	4.44	4.32
Landscapes	3.83	5.50	5.28
Winter	3.07	4.37	4.37

Table 4.4: **Minimum unique entropy of channels in various image corpora.** The minimum unique entropy is extremely variable across the image corpora. Most corpora have two or three bits of unique entropy between the long and medium channels, while the other channel pairs regularly contain more entropy.

4.2.2 Mutual information

Mutual information tells us how much information is shared between channels. Channels could have a high joint entropy, but there may also be high information overlap between the two channels so that a single photoreceptor channel can provide the majority of the information. The mutual information shared between the two channels is computed with

$$I(X; Y) = \sum_{y \in Y} \sum_{x \in X} p(x, y) \log \left(\frac{p(x, y)}{p(x)p(y)} \right), \quad (4.2)$$

where x and y are values that occur within the channels X and Y . The joint probability $p(x, y)$ is the chance that x and y occur simultaneously, while $p(x)$ and $p(y)$ are the independent probabilities that x and y occur within their respective X and Y channels (Yao, 1993).

Table 4.3 shows the mutual information for the image corpora. The mutual information measure shows striking differences among the corpora. The Austin and Winter images have over four bits of shared entropy in the long and medium channels, whereas the Flowers corpus has just over one bit shared between the same two channels. Four bits of shared information is extremely high for two channels that have nine or ten bits of joint entropy. The next section presents how to calculate the unshared entropy within each channel.

4.2.3 Minimum unique entropy

The amount of unique information in each channel can be measured by subtracting the mutual information from the minimum entropy of the two photoreceptor values. The minimum unique entropy for the channels X and Y is

$$M(X, Y) = \min(H(X), H(Y)) - I(X; Y), \quad (4.3)$$

where $H(X)$ and $H(Y)$ are the entropy of the X and Y channels. If the joint entropy was used instead of $\min(H(X), H(Y))$ then the maximum unique entropy would be calculated.

The minimum unique entropy for the corpora are shown in Table 4.4. If two channels have a minimum unique entropy of zero then they are identical. If there is a large minimum unique entropy, then the channels are different and it is possible for a system, biological or artificial, to detect differences. A model that self-organizes using corpora with low unique entropy may have trouble developing neurons that distinguish channels. Each channel in Flowers has over six bits of entropy out of a possible eight and so will be the most likely corpus to train a network that has neurons selective for each channel. The Austin Vicinity corpus, with 2.68 bits of unique information between the long and medium channels, is the most likely corpus to cause a neural network to have trouble organizing neurons that can distinguish between the long and medium channels.

The minimum unique entropy quantifies the observation that LMS channels have more overlap than RGB channels as suggested previously in Figure 2.3 on page 10. Figure 4.2 shows the corpora with the least and most minimum unique entropy (Austin Vicinity and Flowers), and plots them alongside their original RGB channel images. The figure shows the LMS channels have less unique entropy compared to the original RGB channels. The RGB to LMS conversion (described in Appendix A) is reducing the amount of unique channel entropy, and converting it into mutual information.

4.2.4 Redundancy

Mutual information measures the shared entropy between channels, and minimum unique entropy measures how much entropy is not shared. Calculating the redundancy links the these two measures into a single ratio that does not depend upon how many total bits of entropy are in a corpus. If two channels are identical then they are completely redundant. If two channels have no mutual information then they have no redundancy.

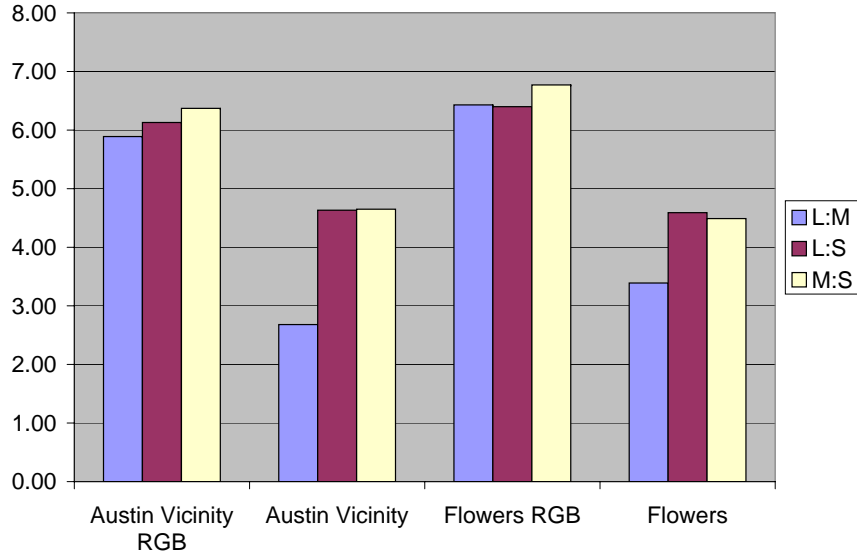


Figure 4.2: **Minimum unique entropy of RGB and LMS natural image sets.** The RGB to LMS conversion described in Appendix A simulates LMS photoreceptors looking at an RGB image. This plot shows that the resulting LMS images have lower minimum unique entropy than the original RGB images. A quarter to a third of the unique channel information has become mutual information because of the overlapping sensitivity functions. Unique entropy is vital to train a network to distinguish between two channels.

The redundancy is calculated with

$$R(X, Y) = \frac{I(X; Y)}{\min(H(X), H(Y))}, \quad (4.4)$$

where X and Y are the long, medium, or short channels of an image. $R(X, Y)$ is 1.0 (100%) when the X and Y channels are completely redundant, and 0.0 (0%) when there is no mutual information between channels (Giasu, 1977).

Table 4.5 shows the redundancy results for the LMS corpora. There is a large variation in redundancy, with sets like Winter having nearly 50% redundancy between channels, yet the most redundant channel in Flowers (long:medium) is no more than 16.5%. The Winter corpus was undistinguished in the other measures but its high redundancy predicts that a LISSOM network may have difficulty self-organizing on it since the channels have more shared entropy than they do unique entropy.

Corpus Name	L:M	L:S	M:S
Austin Arboretum	60.1%	32.0%	33.2%
Austin Vicinity	62.2%	31.2%	31.0%
Flowers	16.5%	13.9%	8.9%
Foliage	47.9%	27.9%	30.0%
Fruit	40.9%	27.0%	28.2%
Kibale, Uganda	56.1%	27.6%	29.6%
Landscapes	48.8%	27.0%	29.5%
Winter	59.9%	43.0%	42.7%

Table 4.5: **Channel redundancy in various image corpora.** Between 25% and 50% of the information found in the photoreceptor channels tend to be redundant information, with the highest redundancy between the long and medium photoreceptors.

4.2.5 Discussion

The results of this section show that out of a possible eight bits each channel value in an LMS triple tends to have a minimum of three to five bits of entropy. These bits represent the information content in the activation of a photoreceptor when presented with an RGB pixel. Each activation value contains an additional two or three bits of entropy, but it is shared with the other channels. RGB channels have only 8% to 16% redundancy, but the LMS photoreceptor activations can be over 50% redundant! As redundancy increases and the minimum unique entropy decreases, it may become harder for the model to detect a unique signal worth becoming selective for. The Flowers and Fruit corpora have high joint entropy and low redundancy, and may be the most useful for training the corpus to develop channel-selective neurons.

4.3 Color distribution

What colors are in the image corpora? To answer this question, representative images from the datasets were converted from RGB to HSV (Hue, Saturation, Value), and histograms of image hues were collected. Results reveal that most images have one or two dominant hues and are extremely biased for those colors. It is hard to find images that contain all of the major colors, and even those that do tend to have uneven distributions. Some hues such as greens and reds occur often, whereas others such as purple are less common.

Figure 4.3 shows three representative images, along with polar plots of the colors found in the images. The left and right images have typical histograms with two or three hues dominating the image. The center image, showing a garden in

Ireland, is uncharacteristic for having many HSV hues visible, and still the histogram is unbalanced.

Across the corpora, the dominant colors found depend upon the subject of the images. The Landscape and Uganda images tend to have more sky visible and thus more blue, but no purple. In contrast, the Flowers corpus contains many purple flowers, but there is less visible sky and the corresponding blue. Landscape images taken in the fall have lots of yellows and reds, whereas spring and summer images have much more green. The winter images are the least chromatic since most of the environment is covered with white snow.

The main conclusion from the histogram study is that not all corpora have the same colors, and therefore the choice of training corpus may affect how the model will self-organize. If a corpus does not have blue, then the modeled cortex may not develop blue-selective regions. It is not yet clear what the most natural colors are in natural scenes (Nascimento and Ferreira, 2002; Parraga, Brelstaff, Troscianko, and Moorhead, 1998; Simoncelli and Olshausen, 2001), but it is reasonable to assume that pictures of nature will capture some of the experiences of the biological eye. However, man-made objects, colors, and environments may have very different characteristics, and are even harder to control and are therefore beyond the scope of this study.


4.4 Principal component of image patches

The previous two sections presented the information content of LMS channels, but they did not discuss the position of the values within the images. What kind of spatial information is contained within the natural images? A great deal of work has already been done in this area using a number of techniques. For example, Independent Component Analysis (ICA) of natural scenes shows a variety of independent components ranging from Gabor-like patches to patterns that resemble receptive fields in the LGN (Lee, Wachtler, and Sejnowski, 2002b; Taylor, Finkel, and Buchsbaum, 2000; van Hateren and Ruderman, 1998; van Hateren and van der Schaaf, 1998). Other studies examine the frequency and type of luminosity edges in natural images and find that edges predominately make gradual transitions (Geisler, Perry, Super, and Gallogly, 2001; Geisler, Thornton, Gallogly, and Perry, 2000). Neural network modelers use grids of neurons trained on natural images and also find that the world contains features that match the receptive fields found in the brain (Bednar, Kelkar, and Miikkulainen, 2002; Miikkulainen et al., 2005; Oja, 1991a; Olshausen and Field, 1996; Simoncelli and Olshausen, 2001).

The goal of this section is to characterize the information in the images at a spatial scale similar to that used by the input regions of the LISSOM model. To achieve this, a Hebbian neuron will be trained on patches of pixels taken from larger




(a) Foliage image




(b) Flower image




(c) Landscape image




(d) Foliage histogram



(e) Flower histogram



(f) Landscape histogram



(g) Hue key

Figure 4.3: **Hue histograms of natural scenes.** Plots (a)–(c) show three color images, each from a different corpus. Plots (d)–(f) are polar plots of the colors found in each image, and (g) is the color key for the polar histograms, starting with red at 0° moving through all the hues and then back to red at 360° . Plot (a) is dominated by orange and blue, whereas (c) is dominated by a broader band of oranges. Plot (b) shows an Irish garden and is unusual for the large number of colors visible, but still the histogram is uneven and far from smooth. These images are representative of the more general finding that natural images contain few colors and have unbalanced hue histograms.

images of natural scenes. The procedure is not new, but the results are, in that they describe the images used in this dissertation, and give the principal component at a scale relevant to the dissertation model. Hebbian learning on a single neuron has been shown to extract the first principal component from the inputs presented to it (Haykin, 1994; Oja, 1982, 1991b). The Hebbian learning rule, implemented as described in Section 2.3.2, has the initial connection weights randomized and the spatially organized receptive field weights found by the network are a direct result of the Hebbian learning.

The rest of this section will first present how the neuron is implemented. Second it will discuss the results of the experiment. And then third, there is a discussion of what the results mean and what they predict about how color training images will affect the LISSOM model.

4.4.1 The simulation architecture

The neuron is implemented in MATLAB with 300 input connections representing activation levels from natural images, and the connection weights adjusted using Hebbian adaptation. There are four RGB images $1,000 \times 1,000 \times 3$ pixels in size from which the training set is generated. Each full image is a natural scene taken with a digital camera and part of the Austin Vicinity corpus (Visual Delights, 1999). The four photographs were chosen randomly from a pool of 75; MATLAB memory constraints made loading more samples difficult on the machine, a 450MHz Pentium-III with 512MB RAM. As part of the Hebbian adaptation algorithm the images were normalized to values between 0 and 1 and then the images were shifted so that the mean of the dataset equals 0. The simulated neuron has a $10 \times 10 \times 3$ receptive field which represents a neuron connected to a small region of the retina where there are all three types of cones. Each training patch is a $10 \times 10 \times 3$ subregion selected at random from one of the full-size images. Figure 4.4 shows three example training patches.

Each training patch is presented to the Hebbian neuron and the connection weight adjustments ($\Delta\omega_i$) are calculated using Equation 2.1. After the weights are updated they are divisively normalized to

$$\sum_{i=1}^{300} \omega_i = 1.0. \quad (4.5)$$

A total of 20,000 random training images are used to adapt the neuron weights. Weight adjustments are scaled so that the adjustments are at first large, and then they exponentially decrease during training. The exponential weight scaling factor




Figure 4.4: **Sample training input patches.** Three sample $10 \times 10 \times 3$ input training patches, extracted from larger images taken of natural scenes. These patches show that at this scale, the objects in the image cannot be identified. Most patches contain a single color, and have gradual luminosity changes like those shown here.

(α in Equation 2.1) begins at e^8 and eventually decreases to 1.0. The core algorithm can be seen in Figure 4.5.

4.4.2 Results

After training, the neuron’s weight matrix contains an orientation- and color-selective receptive field, shown in Figure 4.6. The neuron is maximally selective to stimuli that have an orientation of about 50° from vertical and a spatial frequency of about five pixels.

It is possible to display the entire receptive field as an RGB image as shown in the upper-left plot of Figure 4.6. The luminosity value of a pixel represents the strength of the weight to that input, while the chromatic value specifies which color channel is favored, in this case green. Green weights in the receptive field means that the neuron has strong weights to the green input channels with weaker weights to the red and blue channels. The three contour plots in Figure 4.6 are the 10×10 red, green, and blue channels of the single $10 \times 10 \times 3$ receptive field. These plots show that each channel has the same orientation preference, but the strength of the weights to the green channel are consistently stronger and these alone make this neuron green-selective.

4.4.3 Discussion

The receptive field of a neuron trained with natural image patches becomes primarily orientation selective. In the neuron being studied the structure remained similar to

```

1  ITERATIONS = TRAINING_SIZE * WEIGHT_ADJUSTMENTS;
2  m = ones(RF_SIZE,RF_SIZE,3) * INITIAL_WEIGHTS; % Weights of RF
3  dm = zeros(RF_SIZE,RF_SIZE,3); % Amount to change m
4  norm_dm = dm; % Normalized m
5
6  % Learning Rate Schedule
7  mu = 1 ./ exp(EXP_DECAY_START:-EXP_DECAY_START/ITERATIONS:0);
8
9  ts = TRAINING_SIZE;
10 for i = 1:ITERATIONS
11     d = squeeze(TRAINING(mod(i,ts)+1,:,:,:)); % Current Input
12     c = sum(sum(sum(d .* m))); % Total Neuron Activity
13     dm = mu(i) * (d * c); % Hebb Rule
14     norm_dm = norm_dm + normalize(dm); % Add dm for later
15
16     if mod(i,ts) == 0 % Once each image in the training
17         m = m + dm; % set has been presented, make a
18         m = normalize(m); % normalized weight adjustment.
19         norm_dm = 0; % Do this WEIGHT_ADJUSTMENT times.
20     end
21 end

```

Figure 4.5: MATLAB implementation of the Hebbian neuron. The for-loop shows the weight calculation, adjustment, and normalization step for each iteration. Reprinted so the results can be reproduced if needed.

a neuron trained with monochrome input, but there was a noticeable shift towards green inputs which makes the neuron prefer green stimuli. A different set of source images would make the neuron prefer a different color. The green preference of the neuron can be explained *ex post facto* by studying the training data. The original dataset was shifted so that the mean of the inputs $d_i(x, y, c)$ summed to 0:

$$\sum_{i=1}^{20000} \sum_{x=1}^{10} \sum_{y=1}^{10} \sum_c^{\{R,G,B\}} d_i(x, y, c) = 0. \quad (4.6)$$

If each color channel is summed separately then the hidden structure can be seen. It shows that the green channel has a higher mean:

$$10^2 \sum_{i=1}^{20000} \sum_{x=1}^{10} \sum_{y=1}^{10} d_i(x, y, R) = 1.52, \quad (4.7)$$

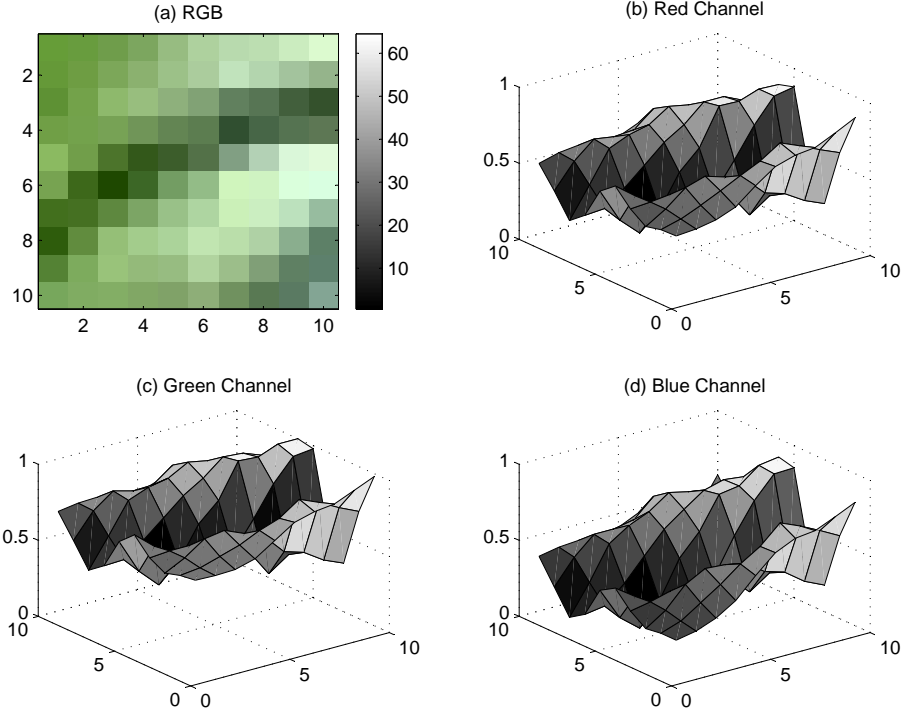


Figure 4.6: **Learned weight function of the Hebbian neuron.** The learned receptive field of the Hebbian neuron is equivalent to the first principal component of the input patches. The shape of the receptive fields are similar for each of the three channels (a–d) even through each weight matrix has a different offset near 0.5. The network learned to have the strongest connection to the green channel in plot (c), noticeable from having the highest overall connection weights and a slight green color in the combined plot (a). This result predicts that the LISSOM model trained on natural images will organize orientation-selective receptive fields with color biases towards the brightest channel.

$$10^2 \sum_{i=1}^{20000} \sum_{x=1}^{10} \sum_{y=1}^{10} d_i(x, y, G) = 3.90, \quad (4.8)$$

$$10^2 \sum_{i=1}^{20000} \sum_{x=1}^{10} \sum_{y=1}^{10} d_i(x, y, B) = -5.31. \quad (4.9)$$

The weight adjustment rule shifted the weights towards green because it was the strongest of the three color channels. However, the luminosity information contained in the input far outweighed the color information. As a result, the weight

matrix extracted primarily orientation information with a slight bias towards the dominant color. It is hard to detect color differences in the three channels by inspecting them visually. That Hebbian learning is able to detect the green channel bias is a remarkable example of the sensitivity of the algorithm.

In conclusion a single Hebbian neuron can become partially selective for color, but is dominated by orientation stimuli. It becomes feature selective by aligning the weight matrix along the strongest color while finding the principal luminosity component in the input data. This study also predicts that the full LISSOM model may become biased towards the strongest colors in the stimuli and the most active photoreceptor channels.

4.5 Color and luminance gradients

Visual inspection of training images suggests that colors change less frequently in images than orientations. This observation can be quantified in terms of a gradient map.

A gradient map measures the change between neighboring pixels in an image. Large changes in the source image creates large values in the gradient map, and areas with no change in the source have a gradient of zero. Because the LGN neurons in LISSOM have Difference of Gaussian receptive fields (which respond to contours and changes in the inputs), a gradient map is an approximate measure of how much LISSOM will be activated by an image.

Figure 4.7a is an HSV (Hue/Saturation/Value) image of some ground ivy. The value component of this image was used to calculate the luminosity gradient in Figure 4.7b, based on the Euclidean distance equation

$$D_{ij} = \sqrt{D_{x,ij}^2 + D_{y,ij}^2}, \quad (4.10)$$

where $D_{x,ij}$ and $D_{y,ij}$ are the differences between the value of pixel (i, j) and its respective vertical and horizontal neighbors (Miikkulainen et al., 2005).

The hue and value components were used to calculate the hue gradient map in Figure 4.7c. Hue gradients are masked by the value pixel components since black areas are prone to noise that can cause large gradients but little to no actual retinal stimulation. Because Hue values in HSV have two dimensions of color (yellow/blue and red/green), the hue distance was calculated with

$$D(\theta_1, \theta_2) = \frac{1}{2} \sqrt{(\cos\theta_1 - \cos\theta_2)^2 + (\sin\theta_1 - \sin\theta_2)^2}, \quad (4.11)$$

where θ_1 and θ_2 are the hue values being compared.




Figure 4.7: **Gradient maps of a single natural image.** Figure (a) is an image of some ground ivy used as a training image for the simulation. (This image is from Olmos and Kingdom, 2004). Figure (b) is a contour gradient map of the luminosity in the image. The amplitude is represented in gray scale from white to black (low to high). The average luminosity gradient is 7.6%. Figure (c) is a contour gradient map of the hue channel from the HSV decomposition of the image. The change in color is represented from white to black (low to high). The average hue gradient is 5.2%, most of which is in dark areas of the image that would cause little retinal activation. Natural images tend to have a larger luminosity gradient compared to the color gradient. Having slower color changes in the training input may be what causes V1 to form isolated color blobs instead of continuous bands.

The average luminosity gradient for Figure 4.7a is 7.6%, and the average hue gradient is 5.2%. This means that on average there are more luminosity contours than there are color contours in this image. The same is true for the majority of the images in the training sets. As suggested by the PCA study in the last section, when the average hue gradient increases or the luminosity gradient decreases, LISSOM maps may begin to be dominated by color selectivity. This result is interesting because it predicts that developing biological cortex requires strong orientation stimuli to form normal cortical maps.

Figure 4.8 shows the average luminosity and hue gradients for each of the training sets presented in this study. The image sets have consistently higher orientation gradients compared to the hue gradients. There are a number of images within the corpora that have greater hue gradient values than luminosity, but they are the




Figure 4.8: **Average gradients of image corpora.** Average orientation and hue gradients for the image corpora, including standard mean error bars. All image sets have a higher average orientation gradient than hue gradient. The large variations between corpora are due to the spatial frequencies found in the images and the quantity of colors. For example, many of the Austin Arboretum images contain many small leaves that are the same shade of green. Larger orientation gradients explain why the principal component study organizes first for orientation, and then for color.

exception.

4.6 Conclusion

These studies reveal five characteristics about the natural image corpora:

First, the three color channels (long, medium, and short) in the retina have highly correlated activations. Second, even with high correlations, most natural images have more unique entropy than they have shared entropy. Third, color is distributed irregularly in natural images, such that in order to represent all the possible colors, a number of images may need to be pooled together. Fourth, oriented edges are the dominant feature of small-scale receptive fields, with color as a secondary feature. And finally, changes in color occur less frequently than changes in luminosity.

There are three predictions that come from the results of this chapter. First, the information theory analysis predicts that the LISSOM model will have difficulty organizing on corpora that have channels with a high shared entropy and low unique

entropy. There is likely to be some minimum required entropy for the model to self-organize correctly. Second, the Hebbian neuron study predicts that the LISSOM model will self-organize towards the channel that has the highest mean luminance, causing a color-selectivity bias in the cortex. And third, the contour gradient study predicts that training a neural network on an image set that has larger hue gradients than orientation gradients will cause the network to prefer colored stimuli over orientation stimuli. Future chapters will study how the features in natural images affect the model's self-organization.

Chapter 5

Modeling self-organizing feature preferences in primary visual cortex

In a self-organizing process based on natural image input, the LISSOM V1 model develops overlapping retinotopically organized color, orientation, and ocular dominance maps. This section analyzes this organization, including novel findings about how color is represented structurally and functionally in V1. Map-level analysis predicts that there are three distinct color-selective structures: red-selective regions, green-selective regions, and blue-selective regions. Color-selective regions are found within ocular dominance stripes, and in the parts of V1 not selective for orientation, as they are in biology. The model further predicts that lateral connections are strong between neurons that have similar color preferences. The model therefore both matches known biological structures and makes further verifiable predictions for future experiments.

5.1 Map level measurements

The LISSOM architecture and training procedure have been described in Chapter 3. After self-organization, the preferences for each neuron in V1 were measured by presenting sine gratings of various colors, orientations, frequencies, and phases. As described below, these measurements show that the model V1 develops topographically organized maps for orientation, color and eye preference.

5.1.1 Orientation-selectivity maps

As in previous monochromatic LISSOM models (Bednar, 2002; Sirosh, Miikkulainen, and Bednar, 1996), most V1 neurons become strongly selective for orientation. Figure 5.1 is a view of V1 with each pixel representing a neuron, colored according to its preferred orientation. The resulting orientation map, which is similar to those found experimentally in animals, containing features such as pinwheels, linear zones, and fractures (Blasdel, 1992b; Landisman and Ts'o, 2002b). The orientation-selectivity map is similar to the monochrome LISSOM simulations, but there are also now patchy unselective regions. As described in the next section, these unselective regions match perfectly with the color blobs shown in Figure 5.1c.

If there are statistically more edges with a certain orientation in the input images, then the LISSOM model will self-organize more neurons to be selective to the dominant orientation (Bednar, 2002). Indeed, a histogram that bins V1 neurons by orientation preference (Figure 5.1d) shows that there are more neurons selective for horizontal orientations compared to other orientations. Since natural scenes statistically have more horizontally and vertically oriented edges than those of any other angle (Switkes, Mayer, and Sloan, 1978), the model is simply following the statistics of the natural scenes that were used as input. This overrepresentation is consistent with experimental studies showing that monkey and ferret V1 has a larger area devoted to the cardinal orientations (Coppola, White, Fitzpatrick, and Purves, 1998; Mansfield, 1974). The horizontal bias in the orientation map is another point of validation for the computational model.

5.1.2 Color maps

Through self-organization, smoothly varying maps of orientation and color preferences emerge in V1. Figure 5.2 shows that color-selective regions in the model V1 are organized into scattered discontinuous patches, or “blobs.” There are three types of blobs: red-selective, green-selective, and blue-selective. Blue-selective neurons also respond to cyan and purple, and both green and red-selective areas also respond to yellow colors. The contour map overlay in Figure 5.1b shows more clearly how neurons selective for color are less selective for orientation, and vice versa. This result is consistent with observations by Livingstone and Hubel (1984).

A unique combination of color-selective blobs activate for each of the six HSV colors (red, yellow, green, cyan, blue, purple), which means that the color of the input can be estimated accurately using only measurements of V1 activation. V1 needs to represent input colors accurately, so that higher cortical regions can construct the full perception of color. Regions selective for color tend to be maximally selective towards one of three different hues as shown in Figure 5.3a, but the neurons in these blobs

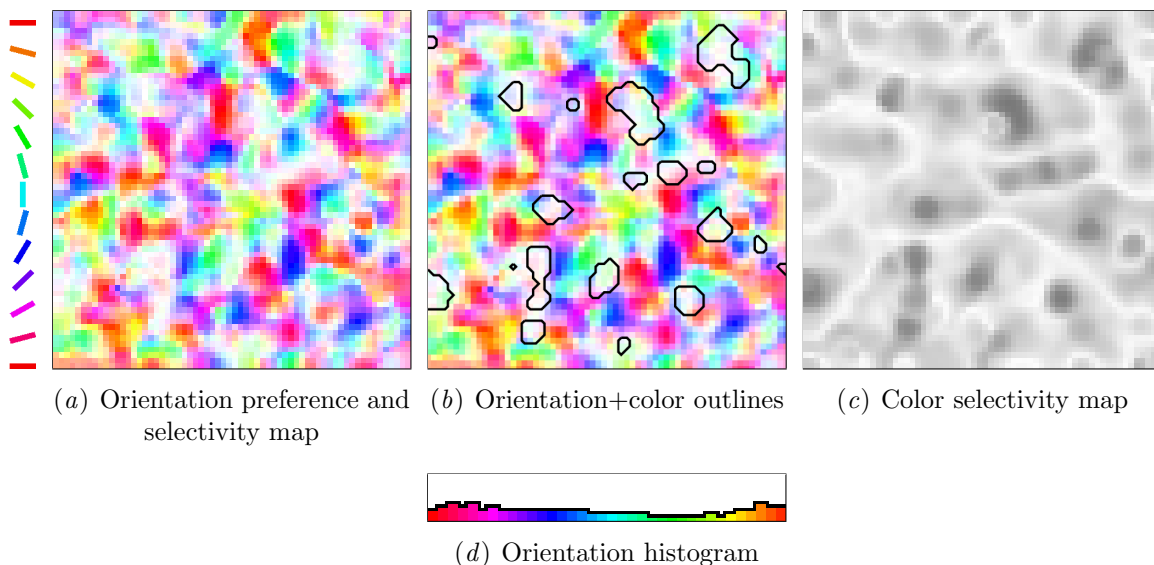


Figure 5.1: **Self-organized orientation and color maps.** Each of the 64×64 neurons is colored with its preferred orientation, according to the key at left. Regions of unselective neurons scattered through the orientation map are colored white and match the regions selective for color, shown as black blobs in (c). Plot (b) outlines the color sensitive regions on the orientation map to make this relationship clear. These orientation and color maps are similar to those found experimentally in the macaque monkey (Blasdel, 1992b; Landisman and Ts'o, 2002b). Plot (d) is a histogram of the orientation preferences for all V1 neurons. Horizontal preferences are more common than diagonal ones, matching the natural scene input statistics, and duplicating how biological orientation maps also prefer the cardinal orientations (Coppola et al., 1998; Mansfield, 1974).

also respond to other colors. Figure 5.2 shows how the cortical responses to different full-field solid colors overlap. The grayscale background of the figure (also shown in Figure 5.1c without the color contours) shows the total overall color selectivity, black being maximally color selective and white unselective. The three blob types seen are red-yellow, green, and cyan-blue-purple. Thus, the model predicts that the biological cortex will have three dominant types of color-selective regions: red-yellow selective, green selective, and cyan-blue-purple selective.

The color-selectivity maps (Figure 5.2 and Figure 5.3a) contain a few regions where red and green blobs touch to create a red-yellow-green hue band. The neurons along the multi-color bands form a continuous transition between colors, mirroring the smoothly continuous colors observed perceptually. The model predicts that smooth transitions will exist between any directly adjacent color-selective regions in V1.

The cortex forms color-selective regions based on the colors that are presented

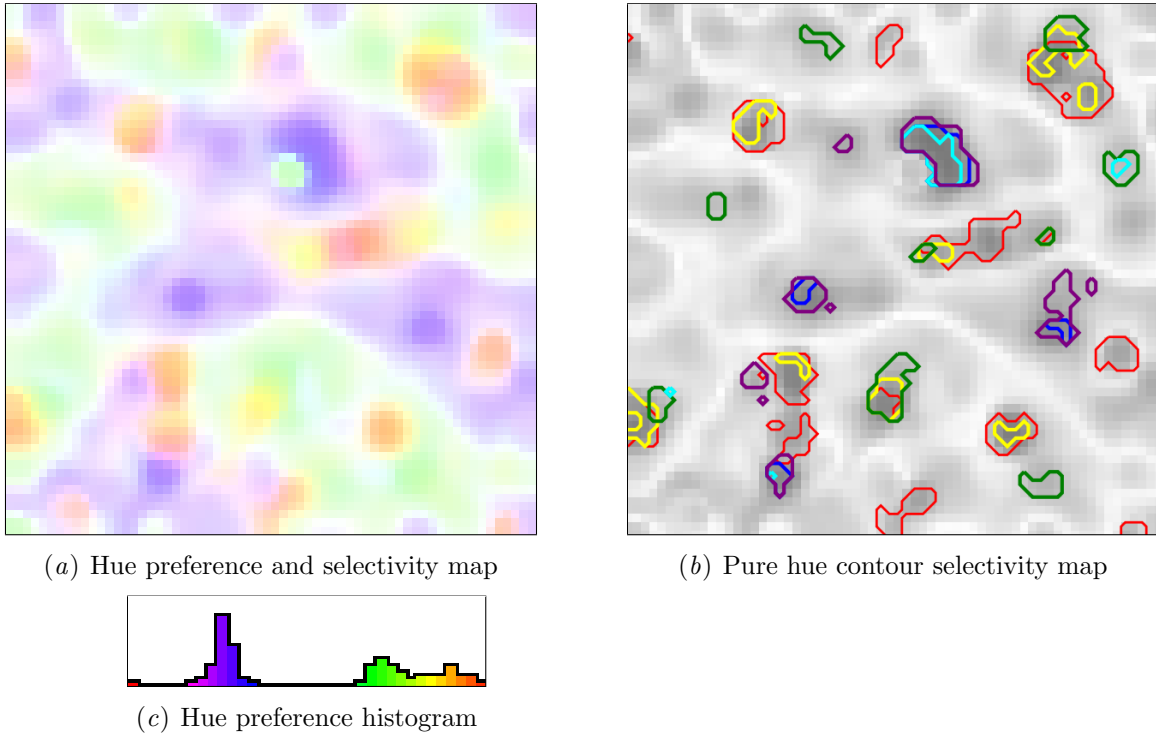


Figure 5.2: **Pure hue color-selectivity.** Plot (a) shows the color preference and selectivity of each neuron, with the histogram of hues in (c). Plot (b) shows the responses of V1 to six HSV hues. For each color, the area that responds is encircled in this plot with a ring drawn in that color. The background grayscale plot shows the strength of selectivity for color in general, as in Figure 5.1c. The three blob types seen here are red-yellow, green, and cyan-blue-purple. Thus the model predicts that biological cortex has these types of color-selective regions. Similarly, in a few cases blobs touch and create hue bands, which are also a prediction of the model.

to the retina during training. However, if a photoreceptor channel has less than three bits of unique entropy per stimulus, as calculated in Section 4.2, the cortex becomes color-blind to this channel. For instance if the cortex sees little short cone activity, it will not become selective for colors dependent on the short cones, as will be discussed further in Chapter 7.

5.1.3 Ocular dominance

As in the macaque (Bartfeld and Grinvald, 1992; Landisman and Ts'o, 2002b), the model color-selective blobs are found within ocular dominance stripes. Figure 5.3a

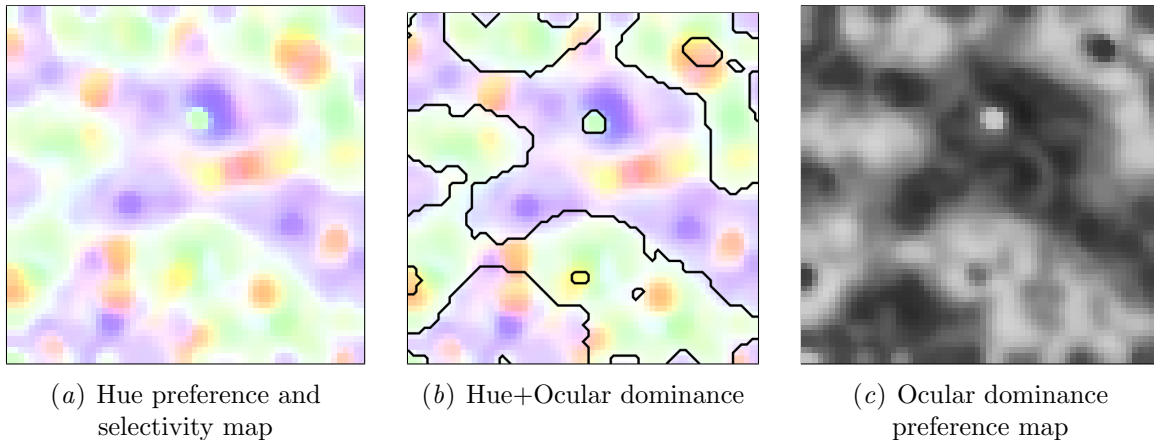


Figure 5.3: **Self-organized color maps and ocular dominance.** Each of the 64×64 neurons in (a) is colored with its preferred color, with color-unselective neurons colored white. In (b), the ocular dominance preference stripes from (c) are outlined on top of the hue-selectivity map to show how the two maps interact. Similar to biology, color blobs are primarily contained within the monocular regions of the simulation (Bartfeld and Grinvald, 1992). These results accurately reproduce the biological organization, further validating the model.

illustrates this relationship. The color of each pixel shows the color preference of the neuron, with color-unselective regions shown in white. Figure 5.3c is the same region of cortex measured for ocular preference, and Figure 5.3b is Figure 5.3a with contour lines drawn along the ocular dominance stripes. Most of the blobs are found within the ocular dominance bands, i.e. most color-selective blobs have a bias toward the left or the right eye. However, there are a few blobs that do not follow this pattern, instead they are split between a left-eye and right-eye preference. Similar cases can be found in biological maps (Landisman and Ts'o, 2002b).

Figure 5.4a compares the ocular dominance maps of the V1 model directly with macaque cortex data from Landisman and Ts'o (2002b). The location of pinwheel centers are marked with yellow dots and color-selective patches are outlined in red. In both maps, regions selective for color tend to be located within the ocular dominance stripes, though the biological and modeled cortex have some color blobs located along the transition between ocular dominance stripes. In both plots the location of pinwheel centers do not appear to be related to either the color-selectivity or ocular dominance maps. Overall, the V1 cortex is consistent with data from the macaque, further validating the model.




Figure 5.4: **Comparing biological and simulated V1.** Plot (a) is Figure 8e from Landisman and Ts'o (2002b, used with permission) showing the ocular dominance map of a 5.3mm by 2.3mm region of macaque monkey V1, with the location of pinwheel centers marked with yellow dots and color-selective patches outlined in red. Figure (b) shows a corresponding plot from the LISSOM model. In both maps, regions selective for color tend to be located within the ocular dominance stripes, whereas the locations of pinwheel centers do not appear to be related to the color-selectivity maps.

5.2 Individual neurons

The previous section reported measurements of the model cortex at the map level. This section analyzes the receptive fields and lateral connectivity of individual neurons in the model V1. Nearly all neurons develop receptive fields selective for color or for orientation, while some neurons are partially selective for both. Most neurons are binocular, but many have a preference for the left or right eye. The subsections below will examine each of these major types of neurons, and make predictions about the afferent and lateral connections of biological neurons.

5.2.1 Color-selective neurons

Each of the three types of color blobs in the model V1 contains a distinct type of color-selective neuron with a specific connectivity pattern to the LGN sheets. Figure 5.5 shows a neuron that has become selective for red stimuli. The large circles in Figure 5.5a are the receptive field projections from each LGN sheet to the color-selective neuron. The top two rows of receptive fields in Figure 5.7a are LGN sheets that receive input from the left eye, and the bottom two rows receive input from the right eye. White pixels within the receptive fields indicate a strong connection to the receptive field, i.e. an area that will excite the neuron if it is activated, while black denotes weak connections that do not excite the neuron. Through the process

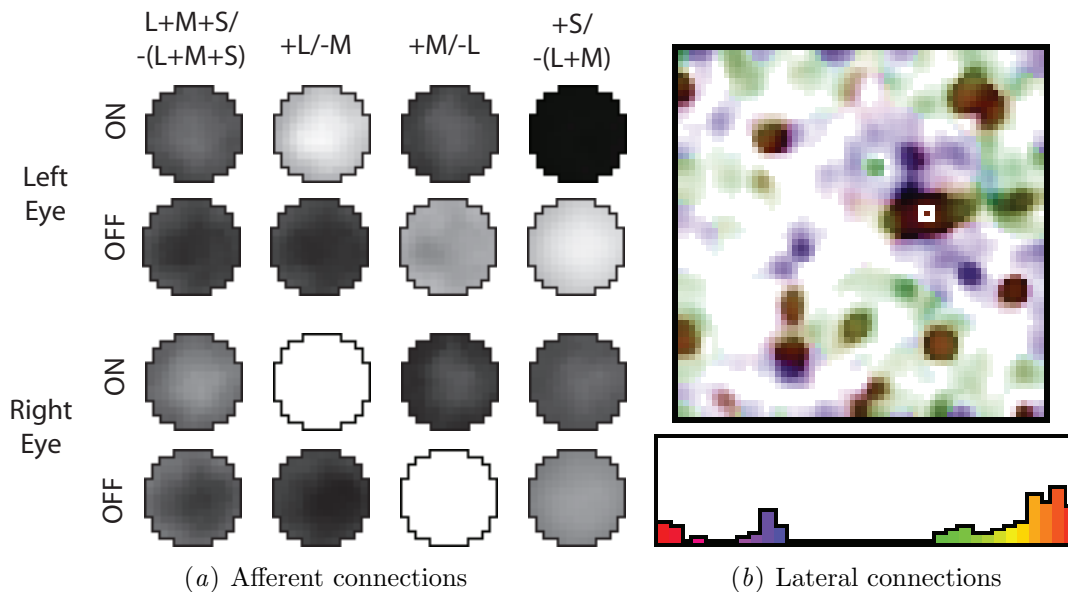


Figure 5.5: **Receptive field of a color-selective neuron (red)**. The large circles in (a) are the receptive field projections from each LGN sheet to the color-selective neuron. White pixels within the receptive fields indicate a strong connection in the receptive field, while black denotes weak connections. The top two rows contain receptive fields to left eye LGN sheets and the bottom two contain the LGN sheets connected with the right eye. In (b), the lateral connections of this neuron are shown, with the neuron's position marked with a small white box. Strong lateral connections in the map are dark, weak connections are gray, and the absence of a connection is shown in white. Lateral connections are colored according to the hue preference of the source. The histogram at the bottom of (b) bins the lateral connections by the hue preference of the connected neurons. The neuron has strong connections to the L/-M, -M/+L, and -S/+L+M sheets and forms lateral connections primarily with other red-selective neurons. The model predicts that color-selective neurons will have connectivity patterns similar to these.

of self-organization, the neuron has near-zero connection weights to all LGN sheets except the L/-M, -M/L, and -S/+L+M sheets. Though the neuron is most selective for red, it will also respond to yellow light, which activates those LGN sheets as well. Intuitively, the connected receptive fields can be thought of as red but not green (long ON/medium OFF), not green but red (medium OFF/long ON) and not blue but anything else (short OFF/long ON + medium ON). These three receptive fields of the neuron are fully connected with no fine structure, so the neuron will be stimulated by any activity within the receptive field. Thus these neurons are not spatially center-

surround, and respond best to large patches of the preferred color. Red-selective neurons are maximally activated by the red-orange hue with which the red-selective blobs are colored in Figure 5.3a. The specific color preference of red-selective neurons is a prediction of the model.

Neurons of similar orientation or color selectivity connect laterally to other areas with the same functional preference. For example, neurons with a preference for vertical orientations tend to connect to other areas that prefer vertical orientations. Likewise, neurons selective for red connect to other regions selective for red; neurons selective for green connect to other regions selective for green; neurons selective for blue connect to regions selective for blue. These connections reflect the patterns of correlations during self-organization.

In Figure 5.3b, the lateral connections of the red-selective neuron (which is marked with a small white box) are shown. Strong lateral connections in the map are dark, weak connections are gray, and the absence of a connection is shown in white. The darkest blob consists of short-range connections from nearby neurons. The lighter blobs consist of long-range connections from distant patches of neurons. Lateral connections are colored according to the hue preference of the source. The sample neuron primarily makes strong connections with other neurons selective for red stimuli, but in addition there are a number of connections to weakly color-selective cells that have a green or blue hue preference but are not selective enough to form a color blob. These weak connections show up in the histograms even though most strong connections are to the highly color-selective regions. The histogram at the bottom of Figure 5.3b bins the lateral connections by the hue preference of the connected neurons.

This primarily red-selective cell is one of the three types of color-selective neurons; green-selective and blue-selective cells are similar but have connections to different sets of LGN sheets. These types of neurons arise automatically through the self-organizing process, driven by the learning rules, LGN processing, and training images. The lateral connectivity of the regions selective for color is not known, and these results constitute a prediction of the model.

5.2.2 Orientation-selective neurons

Neurons outside of color-selective blobs become selective for orientation, while cells that are on the edge of color-selective blobs tend to be partially selective for both color and orientation. Figure 5.6 shows a neuron that is strongly selective for an orientation of about 50° . Oriented stimuli are most effective in driving this neuron, but the strong connections to the -S/+L+M sheet connected to the left eye cause this neuron to have a slight preference for yellow stimuli, though it also responds robustly to other colors and to monochrome stimuli. The receptive fields of such neurons are




Figure 5.6: **Receptive field of an orientation-selective neuron.** This figure follows the same organization as described in Figure 5.5. Monochrome stimuli with an orientation of about 50° are most effective in driving this neuron, however the receptive field to the $-S/+L+M$ sheet connected to the left eye, creates a slight bias towards yellow stimuli. In plot (b) the color of the weights correspond to the orientation preference of the connecting neurons (shown with the orientation key at right), and not their color preference. The histogram at the bottom of (b) bins the lateral connections by the orientation preference. In general, orientation-selective neurons contact other orientation-selective neurons with similar preferences, further validating the model (Sincich and Blasdel, 2001).

similar to those found in biology (Conway, 2001; DeAngelis, Ohzawa, and Freeman, 1995; Landisman and Ts'o, 2002b).

Within the model V1, neurons selective for similar orientations are laterally connected to each other, reflecting long-term patterns of correlation. In Figure 5.6b, the dark regions represent the self-organized lateral weights from each neuron in the V1 sheet to the neuron marked by the white box. The darkest blob consists of short-range connections from nearby neurons, which have a variety of orientation preferences. The lighter blobs consist of long-range connections from distant patches of neurons, which prefer similar orientations. In general, orientation-selective neurons contact many other patches of orientation-selective neurons, but do not connect to color-selective neurons. The histogram at the bottom of Figure 5.6b bins the lateral connections by their orientation preference. The color of the bin corresponds to the




Figure 5.7: **Receptive field of an eye-selective neuron (blue color, right eye).** This figure follows the same organization as described in Figure 5.5. The ocular dominance of the neuron is the result of stronger connections to LGN regions that receive input from the right eye. The histogram at the bottom of plot (b) bins the lateral connections by the hue preference of the connected neurons, and shows that most strong connections are to other blue-selective regions, with a few weak connections to green-selective neurons. These eye-selective neurons are a validation of the model (Löwel, 1994; Löwel and Singer, 1992).

orientation preference of the connecting neuron (with the orientation key at right), and not its color preference. The histogram reveals that the neuron is biased towards connecting to other neurons with a similar orientation preference. These results match previous experimental results in the tree shrew (Bosking et al., 1997) and the macaque (Sincich and Blasdel, 2001), further validating the model.

5.2.3 Eye-selective neurons

Most neurons in the model V1 are binocular, with strong connections to both eyes. Neurons near the centers of ocular dominance stripes, however, are often nearly monocular, with stronger connections to one of the eyes. Many such neurons are also selective for color or orientation. Figure 5.7 shows an example of a typical eye-selective neuron and all of the connections to the neuron. This neuron responds best

to a blue patch in the right eye because of strong connections to the +L/-M, +S/-(L+M), and -M/+L sheets. The ocular dominance of this neuron is the result of stronger connections to LGN regions that receive input from the right eye. Other eye-selective neurons will instead be selective for other colors or specific orientations.

Figure 5.7b shows the lateral connections of this neuron, colored according to the hue preference of the target. The histogram at the bottom of Figure 5.7b bins the lateral connections by the hue preference of the connected neurons, and shows that the presented neuron primarily connects with other neurons selective for blue stimuli, and to a lesser extent green-selective and red-selective neurons. The formation of eye-selective neurons is a validation of the model. Meanwhile, the result that blue-selective neurons connect strongly to other blue-selective neurons, and rarely connect to red and green neurons, is a prediction of the model.

5.3 Discussion

The results with the LISSOM model of V1 show that a biologically plausible topographic organization for color, orientation, and eye preference can develop from simple learning rules and exposure to natural scenes with oriented edges and colored patterns. The details of this organization in the model constitute predictions for the first stages of color processing in the cortex. For instance, the model predicts that color-selective neurons will have lateral connections primarily to other color-selective neurons, and that color blobs in V1 will each respond best to one of the three photoreceptor cone types (long/red, medium/green, or short/blue).

The network is robust in that a wide variation in parameters and inputs will produce a similar organization. However, some abnormal data sets will create interesting results. For example, if the input images are heavily weighted towards a certain hue, such as green, then more neurons and blobs selective for that hue will develop. Similarly, if there are no training images that contain hues activating the short-wavelength cones (blues and purples), then the final network will not contain regions responsive to those hues. Later in Chapter 7, we will study why V1 organizes feature maps in this way.

Though blue-selective blobs do respond to magenta (red plus blue) stimuli, no V1 neurons preferred magenta over all other hues, even though some V1 neurons found in macaque do (Xiao et al., 2007). For some artificially modified training image sets, the modeled V1 does create neurons preferring each of the possible hues, including magenta. These experiments and their results will be discussed in Chapter 7.

Nearly all neurons are selective for color, or for orientation, and a few are selective for both. Moreover, color-selective neurons organize into spatially segregated blobs across the cortical surface, while orientation-selective neurons form large,

smoothly varying areas. These results hold over a wide range of parameter values and image datasets, and reflect differences in how orientation and color information is organized in natural images. Specifically, orientation varies quickly (over a short spatial scale), while color changes relatively slowly, with few abrupt edges. The result is that the network organizes primarily for orientation, but also contains some neurons sensitive to color information, at regular intervals.

5.4 Conclusion

The LISSOM model of color selectivity, orientation preference, and ocular dominance preferences replicates the known data on the organization of color-selective neurons in V1, and provides a detailed explanation for how this selectivity can develop through simple learning rules. The model provides concrete and novel predictions for future experiments about what lateral connectivity patterns will be found, and about how the orientation, ocular dominance, and color maps interact. The model thus represents both an embodiment of our current knowledge about the organization of color processing in the early visual system, and a platform for testing that knowledge and generating predictions to guide experiments. In this way, the model will help us better understand the physiological basis for color perception and visual perception in general.

By using the V1 cortex as an input layer, it is now possible to model the V2 cortical region, which will be the topic of the next chapter.

Chapter 6

Modeling self-organizing feature preferences in secondary visual cortex

The modeled V1 presented in the previous chapter makes it possible to study how higher levels of the visual hierarchy self-organize to represent feature preferences. In particular, this chapter focuses on V2 and presents the results of a modeled V2 cortex that receives input from V1 and has neurons with larger visual fields compared to V1 neurons.

V2 self-organizes using the same principles as V1. The V2 neurons self-organize their receptive fields and lateral connections like the V1 neurons, but V2 neurons are also able to integrate over a larger area of the visual field. Like V1, the LISSOM model of V2 self-organizes three types of color-selective blobs, each with a different color preference for red, green, or blue hues. Similar to V1, some V2 neurons prefer secondary colors and are located along the regions where color blobs touch. The V2 cortical map and the individual neurons contained within it model the thin-stripe and inter-stripe regions of biological V2. Interestingly, the model does not organize into a striped pattern, even though neurons within the artificial V2 match biological V2. Theories as to why stripes did not appear in the model will be addressed later in Section 6.3.

6.1 Map level measurements

V2 is trained with binocular training stimuli after V1 has fully developed. The LISSOM architecture and training procedure has been described in Chapter 3. Through self-organization, maps of orientation and color preferences emerged in V2. The pref-

erences for each neuron in V2 were measured by presenting sine gratings of various colors, orientations, frequencies, and phases. As described below, these measurements show that the model V2 develops topographically organized maps for orientation and color preference.

6.1.1 Orientation maps

Most V2 neurons become selective for orientation, though less so than V1 neurons (Figure 6.1). The orientation-preference map has linear zones, pinwheels, and fractures like V1 maps; however the V2 maps are more blob-like in appearance than the V1 maps, which tend to be more even and continuous. As in V1, there are also large areas of orientation unselective regions within the maps, and they correspond to color selective regions. Figure 6.1b shows how neurons selective for color are less selective for orientation, and vice versa. It is validation of the model that color-selective areas within V2 are less selective for orientation (Levitt et al., 1994; Shipp and Zeki, 2002a). As in V1, the V2 orientation-preference histogram shows a slight bias for horizontal orientations.

The orientation-preference maps in V2 self-organize as though performing a low-pass filter of V1 maps. Large preference patches from V1 are mirrored at the V2 level, but small regions of selectivity in V1 for non-dominant orientations do not appear. Also, any biases for particular orientations in V1 maps are magnified within V2 maps so that V2 orientation-preference histograms are less balanced than the corresponding V1 orientation-preference histograms. Similarly, it was found that the total number of V2 neurons selective for orientation directly depended on the total number of V1 neurons selective for orientation. Thus the model predicts that uneven distributions of feature-selective V1 cells will be magnified in V2.

6.1.2 Color maps

Figure 6.1c and Figure 6.2a show that color-selective regions in the model V2 are organized into scattered discontinuous patches. As in V1, there are three dominant types: red-selective, green-selective, and blue-selective. Regions selective for color tend to be maximally selective towards one particular hue but each of the three blob types will also respond to other hues. Red and green blobs also respond to yellow light, and blue blobs also react to cyan and purple light. To illustrate this property systematically (as in Figure 5.2), Figure 6.2b shows overlapping selectivity of the cortex when presented with different full-field solid hues. The grayscale background of the plot, shown in color in Figure 6.2a, displays the total overall hue selectivity, with black being maximally color selective and white being unselective. Thus the model

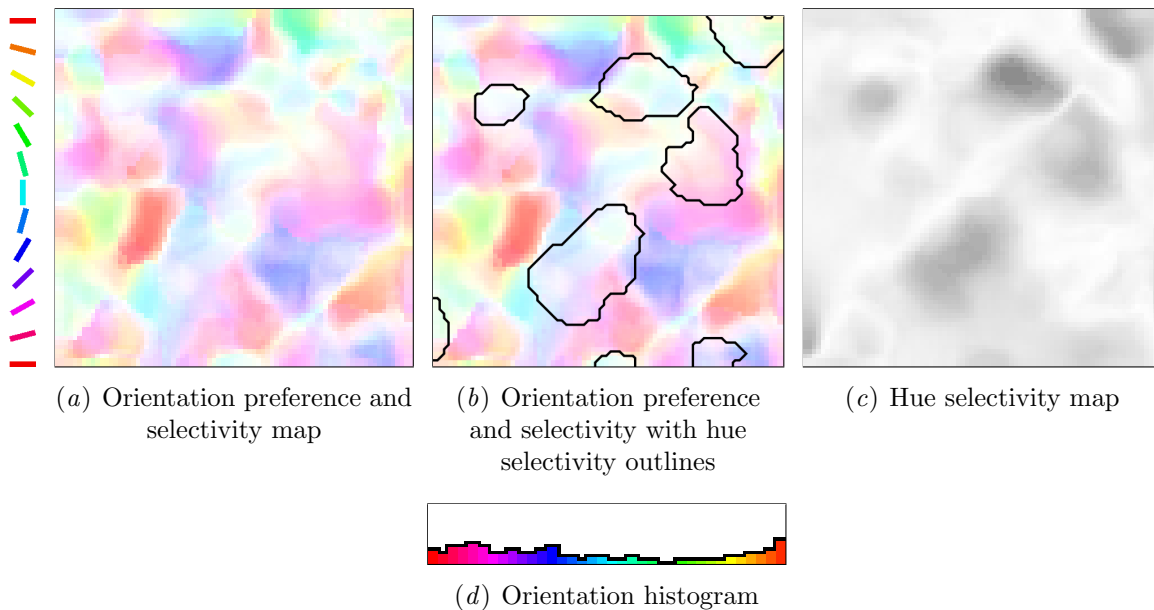



Figure 6.1: **V2 cortex trained using natural color images.** Each of the 80×80 neurons is colored with its preferred orientation, according to the key at left. There are regions of unselective neurons scattered through the orientation map; these are colored white. The orientation unselective regions match regions selective for color, shown as black blobs in (c). Plot (b) outlines the color-sensitive regions on the orientation map to make this relationship clear. Plot (d) is a histogram of the orientation preferences for all V2 neurons. It shows a slight bias for horizontal orientations, which reflects a similar bias found in V1.

matches biological cortex by having color-selective neurons, while also predicting that V2 neurons organize three major types of color-selective regions.


There is one region in the lower-center of Figure 6.2b where a red and green blob touch to create a red-yellow-green hue band. The neurons along this multi-color band form a continuous transition between colors. The model thus predicts that smooth transitions will exist between directly adjacent color-selective regions. The simulation presented here has one transition, but other simulations with artificially modified colors have a varied selection of color transitions (as will be discussed in Chapter 7), suggesting that such organization is a reliable result of the self-organizing process.

6.1.3 Ocular dominance maps

Unlike in V1, eye preference (ocular dominance) maps do not appear in biological V2 cortex (Ts'o et al., 1990). When trained with binocular inputs, where the left-



(a) Hue preference and selectivity map



(b) Pure-hue selectivity contours on hue selectivity map



(c) Hue histogram

Figure 6.2: **V2 hue preference and selectivity maps.** Each of the neurons in plot (a) is colored with its preferred hue with hue-unselective neurons marked in white. The grayscale background in plot (b) represents how selective the neurons are for their preferred hue. Regions selective for color are maximally selective towards the hue shown in plot (a), but the neurons also respond to other hues. To measure these additional responses, full-field stimuli were presented using six HSV hues. The area that responds in each case is outlined in plot (b) with the presented hue. Plot (c) is a histogram of all color selective pixels in plot (a) binned by color preference. The model predicts that biological cortex has distinct types of color-selective regions that blend when they touch other color-selective regions, as shown here where a red and a green blob form a red-yellow-green band.

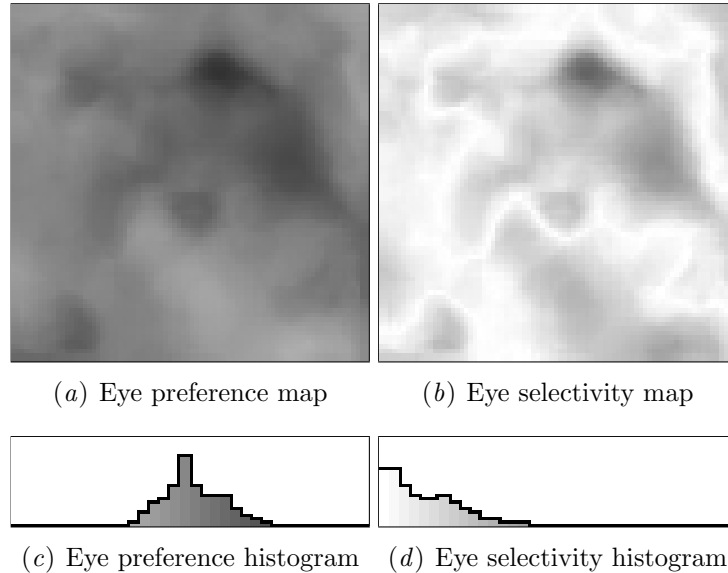


Figure 6.3: **V2 eye preference and selectivity maps.** The V2 cortex does not contain a structured ocular dominance map, just as biological cortex does not (Ts'o et al., 1990). Plot (a) shows each neuron coded by eye preference. White represents a neuron preferring the left eye, black preferring the right eye, and gray preferring both eyes equally. Plot (b) shows the selectivity for eye preference: Few neurons are selective. Plots (c) and (d) bin the neurons of the V2 map, and reinforce the conclusion that V2 neurons are not selective for the eye of origin. These results are similar to those in biology (Ts'o et al., 1990).

eye and right-eye stimulus is identical, the LISSOM model also does not organize a V2 ocular dominance map, as shown in Figure 6.3. The map was measured in the same manner as V1 ocular dominance maps, and yet no structured eye preference or selectivity is seen. All neurons have an eye selectivity near zero and all neurons have preferences close to 0.5. These values mean that V2 neurons do not respond strongly to eye-selective stimuli, i.e. they are binocular.

Interestingly, it is possible to force the model V2 to develop ocular dominance maps by presenting monocular training stimuli, that is, where the left-eye and right-eye inputs have different strengths. It is also possible to adjust the maps by changing the inputs from binocular to monocular during training. Such results raise an important question: Does biological V2 develop during or after V1 self-organization? This issue will be addressed in Section 6.3.

6.2 Individual neurons

The previous section reported measurements of the cortex at the map level. This section analyzes the receptive fields and lateral connectivity of individual neurons within V2. A key difference between V2 and V1 neurons is that V2 neurons receive input from twice the visual field radius of a V1 neuron, measured by back-projecting the afferent connections of the neurons. The V2 receptive field radius is hardwired into the network architecture but the final connection weights self-organize during network training. Nearly all neurons become selective for color or orientation, while a few neurons are selective for both. Meanwhile, all V2 neurons are binocular such that they are stimulated by activation in their receptive field regardless of which eye sees the input. These results are validations of the model. The subsections below will examine the two dominant types of cells: color-selective neurons and orientation-selective neurons.

6.2.1 Color-selective neurons

On average, V2 neurons are less selective for color than V1 neurons. A likely reason is that larger receptive field allows V2 neurons to integrate over a wider area, but weakens the maximum response of the neuron to any one location.

Each of the three types of color blobs found in the model V2 contains a distinct type of color-selective neuron with a unique connectivity pattern. Since each V2 neuron has a single spatially organized receptive field projecting from V1, Figure 6.4 can show the receptive fields from multiple neurons within a single grid. The receptive field of every fifth V2 neuron is shown, arranged according to the neuron's position in the cortex. A strong connection between the V2 neuron and a V1 neuron is represented by a dark spot, whereas a weak connection is white. The color of the connection weight is the color preference of the corresponding V1 neuron, and the saturation represents the color selectivity of that neuron. Strong connections to color-selective V1 neurons create color-selective V2 neurons. For example, blue-selective V2 neurons are maximally activated by the blue-selective V1 neurons within their receptive fields. Similarly, if a V2 neuron's receptive field is dominated by connections to red-selective V1 cells then the V2 neuron is red-selective. Neurons at the edge of the color-selective regions become partially selective for color as well as orientation.

Neurons with similar color and orientation preferences connect laterally to other neurons with similar preferences. These connections represent the patterns of correlation during learning. The lateral connections of three representative color-selective neurons are displayed in Figure 6.5. Each neuron has strong lateral connections to other color-selective neurons in the same color blob as itself, and also has weaker long-range connections to other color blobs with similar chromatic pref-




Figure 6.4: **Afferent receptive fields of V2 neurons encoded for hue preference.** The receptive field of every fifth neuron in V2 is shown. Connections within the plot are colored according to the hue preference of the source V1 neuron. Weak connections are white, and strong connections are plotted using dark or bright colors. The greater the saturation of the color, the greater the selectivity of the V1 neuron for that hue. These types of receptive fields arise automatically through the self-organization process, driven by the learning rules, LGN processing, and training images, further validating the model.

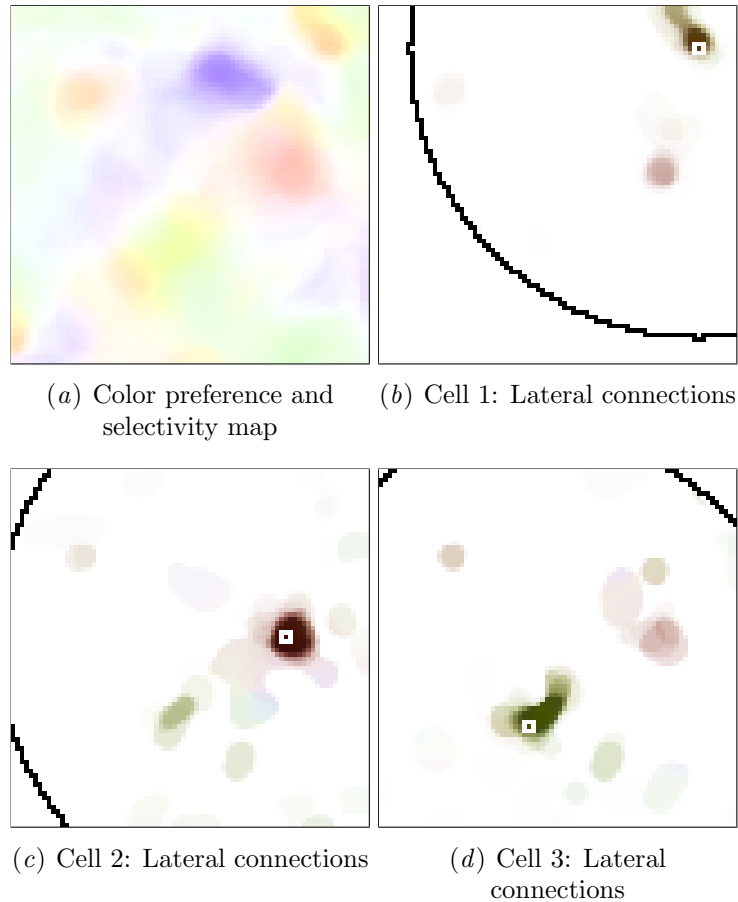


Figure 6.5: **Lateral connections of three V2 color-selective neurons.** Plot (a) shows the hue map of V2 with color-selective neurons colored according to their preference and selectivity. Plots (b)–(d) show the lateral inhibitory connections for three representative color-selective neurons. Each neuron is marked by a small white box. These neurons have strong lateral connections to neurons within the same color blob and also form weaker connections to other color blobs that have similar spectral preferences. These lateral connection structures are predictions of the model.

ferences: Red-selective neurons connect to other red-selective neurons, green-selective neurons connect to green-selective neurons, and blue-selective neurons connect to blue-selective neurons. The model thus predicts that color-selective neurons in biology will form lateral connections to other neurons that have similar chromatic preferences.



Figure 6.6: **Afferent receptive fields of V2 neurons encoded for orientation preference.** As in Figure 6.4, the plot shows the receptive field of every fifth neuron in V2, with strong connection weights plotted using dark or bright colors, and weak connections plotted in white. The strong weights are colored according to the orientation preferences of the V1 neurons, using the color key at right. The greater the saturation of the color, the greater the selectivity of the V1 neuron for that orientation. These types of neurons arise automatically through the self-organizing process and further validate the model.

6.2.2 Orientation-selective neurons

Neurons outside of color-selective blobs are selective for orientation. Figure 6.6 shows the same grid of V2 receptive fields as those shown in Figure 6.4, the only difference being that in Figure 6.6 the V1 neurons are color coded according to their orientation preference. Each orientation-selective V2 neuron tends to form an irregularly shaped receptive field and connect to a single type of orientation-selective V1 neuron. For example, V2 neurons selective for horizontal orientations have receptive fields dominated by horizontal-selective V1 neurons that are colored red. However, there are a few neurons (e.g. neuron (4, 5) from top-left), whose receptive field appears circular and achromatic in Figure 6.6. These are color-selective V2 neurons that connect to color-selective V1 neurons and are a validation of the model.

6.3 Discussion

The simulation shows that LISSOM can create orientation-selective and color-selective V2 neurons. Orientation-selective neurons prefer edges of light with no regard to the color of the stimuli. Meanwhile, color-selective neurons are more responsive to colored stimuli and less responsive to achromatic inputs. These modeled neurons mimic the neurons found in the thin- and inter-stripe regions of biological V2.

However, the LISSOM V2 is conspicuously missing stripes. In biology, thin-stripe regions containing color-selective cells are seen when biological V2 is stained for cytochrome oxidase (Shipp and Zeki, 2002a), as previously reviewed in Chapter 2. Do the missing stripes invalidate the model? Not necessarily. The model has constructed V2 neurons selective for color and orientation just like biological cortex, and these neurons do form connections to V1 neurons just as biological neurons do (Sincich and Horton, 2005b). Individual neurons look like those found in thin-stripe and inter-stripe regions, the only difference is that in the model the color-selective blobs do not form stripes.

To duplicate biological cortex stripes, the color-selective blobs in the model need to align themselves into a columnar organization. It is possible that such alignment did not happen because some necessary feature type was not presented to the retinas. For example, disparity-selective neurons are found in thick stripes, and the model was never presented with disparity stimuli. Perhaps disparity selectivity must be included in the model before stripes will form within V2. An alternative hypothesis is that input from the pulvinar is critical for stripes. Both the thin and thick stripes are known to receive projections from the pulvinar but the role of these connections is still unclear (Sincich and Horton, 2005a).

The model also raises an interesting question about the development of V2.

In what phase of usual system development does the V2 color-selectivity map form? Within the model, V2 can either self-organize simultaneous to the V1 map, or it can self-organize after the V1 map has already formed. These two methods generate similar, but not identical, results. Training V1 before V2 keeps ocular dominance maps from forming in V2, but also causes V2 to develop weaker selectivity maps and broader color-preference blobs. Therefore it is a prediction of the model that V1 ocular dominance maps must self-organize before V2, perhaps driven by spontaneous retinal activity patterns that occur before V2 neurons develop selectivity (Wong, 1999).

6.4 Conclusion

In conclusion, a LISSOM model of V2 was presented that uses a simulated V1 cortex as input and self-organizes a number of biologically plausible features such as color-selective and orientation-selective neurons. The model also makes explicit predictions about the lateral connectivity of neurons. It can therefore serve as a platform for future studies of V2 function, allowing us to understand and reconcile the results from different labs, and perhaps eventually unifying them into one consistent computational model.

The V2 simulation presented in this chapter created a hue-preference map that contained many but not all possible hues. The next chapter will present simulations that contain all hues in the hue-preference map, as well discuss why some colors appear less often than others. In addition the next chapter will also analyze other variations of the model, thereby characterizing the limits of the self-organizing process.

Chapter 7

Experimental Analysis

The previous two chapters presented the complete two-eye trichromatic LISSOM model. It is the most complete retina-to-V2 map model to date and provides several insights into how color is represented in the visual system. This chapter analyzes the model and tests the limits of color selectivity in the LISSOM cortex, to help understand why color maps organize differently from orientation maps and ocular dominance maps. Similarly, this chapter studies what conditions are necessary to develop realistic results, because some types of training inputs self-organize biologically realistic maps, and others do not.

The first section describes how variations in training stimuli affect the model, and why color maps form patchy regions of selectivity. The second section shows that simulations trained with natural images are similar to simulations trained with randomly colored images, suggesting that color-preference maps are rarely affected by the strong color biases in natural images. However, training a simulation with natural images is not sufficient to get the most realistic V1 and V2 maps. The third section shows what is necessary for LISSOM simulations to form V1 and V2 preference maps that include all possible colors, closely matching biological data.

7.1 Sensitivity to image features

LISSOM is a robust model, and will create maps with similar features with a wide range of training images and simulation parameters. However, a balance of ocular dominance, orientation, and color information is necessary for the cortex to develop normally. Too little color and no color-selective regions will form; too much color and the cortex will be overwhelmed by color-selective regions to the detriment of the other feature maps. Color-selective regions will only form into biologically realistic blobs if there is the right amount of color information. Measuring the correct balance

can be described in terms of r^2 correlations and gradient maps like those presented in Chapter 4.

7.1.1 Channel correlation

Pearson's r^2 correlation measures the correlation between two channels, while ignoring the spatial location and the average luminosity of the individual pixels. The three photoreceptor types lead to three r^2 values: long:medium, long:short, and medium:short. In all cases where each of these three r^2 values for a training set was less than 0.99, the network organized three types of color blobs. Conversely, whenever the correlation between any two channels was greater than 0.99, the V1 and V2 regions developed a single neuron type sensitive to both channels. In such cases, the network was unable to differentiate between the two colors, thereby exhibiting cortical color blindness.

The correlation between the long and medium channels was almost always higher than the others because their wavelength sensitivity functions have such a high degree of overlap. The r^2 test thus reveals that a low channel correlation in the retinal stimuli is vital for V1 and V2 cortex to develop realistically, and suggests that cortically caused color blindness can arise if the two channels are not distinctive enough. Color blindness is a prediction of the model and could be tested on developing animals by raising trichromatic animals in environments with only red and green or monochromatic stimuli. Animals raised in this environment would not have regular trichromatic vision, unless the cortex is able to rapidly reorganize once normal stimuli are presented to the animal as an adult.

Conversely, if the color channels in training stimuli have too low a correlation, then the V1 and V2 color-selective regions become larger, merge, and no longer look like spatially separated blobs. In such cases, color-selectivity maps instead show large bands and clusters, and V1 and V2 have more color-selective neurons than they do orientation-selective neurons. When the number of neurons that prefer oriented stimuli decreases, the orientation-selectivity map becomes poorly organized and does not have as many pinwheels and linear zones as are normally found.

7.1.2 Feature map interaction

The last subsection explained that channel correlations change the preferences and shape of color-selective blobs. Similarly, the ratio of color-selective to orientation-selective cells in V1 and V2 also depends upon the hue gradients and orientation gradients of the training images. Some training corpora regularly create more and larger color-selective regions and smaller orientation-selective regions than the simulation presented in Chapters 5 and 6, even though the r^2 values were within normal

ranges. The images in these corpora were consistently found to have large hue gradients relative to their orientation gradients. Likewise, it was found that if an image set has a small average color gradient then a network trained on the set will have few or no color-selective regions. These results predict that V1 and V2 must have oriented and colored stimuli properly balanced, with orientation stimuli stronger than colored stimuli, for normal V1 and V2 development. Such a prediction can be tested experimentally by modifying the environment of developing animals.

The network training stimuli also affect the location of color-selective blobs in V1 ocular dominance maps (V2 does not have a structured OD map). In the Chapter 5 V1, the location of color blobs in the color-selectivity maps matched the location of the ocular dominance stripes in the eye preference map. However, this organization can be altered by selecting different types of training inputs. If the training images have a low correlation between channels (which increases the color), the color blobs will often overlap the ocular dominance stripe boundaries. Similarly, decreasing the correlation between the stimuli presented to the left and right eye will also cause color blobs to overlap the ocular dominance boundaries. This result suggests that either the ocular dominance or the color map can become dominant, depending on the input statistics, just as previously found for orientation and motion maps (Miikkulainen et al., 2005).

The model thus predicts that a species or individuals that pay more attention to colored stimuli, or raised in more brightly colored environments, or experience fewer salient oriented contours, will have color blobs that extend across ocular dominance stripe boundaries. The prediction could be tested by raising animals in environments with strong color and then measuring the location of color blobs in V1.

7.2 Random color distributions

The quantity and types of colors vary greatly in natural images. Since natural images were used as input in Chapters 5 and 6, the results were affected by the statistical characteristics of the images. Such an effect is desirable, since the biological cortex is also affected by the natural environment. However it would be useful to characterize the color-selective properties of the model with no externally introduced biases. The following simulation reveals that natural colors in the world actually lead to very similar organization compared to randomly generated colors.

7.2.1 Training stimuli

To create an unbiased simulation it is necessary that every variable remain constant while the color of each individual stimulus is randomized, eliminating any individual

color dominance. It is also crucial that the saturation of the hues is not accidentally increased. A training set with unbiased colors can be derived from a set of natural images if the original biased colors in the images hues are randomized just prior to presenting an image patch to the retina. For this simulation, 20,000 input stimuli were used, sampled from 23 images in the McGill image database (Olmos and Kingdom, 2004), since images from the McGill corpus were also used for the results presented in Chapters 5 and 6. For each of the 20,000 presentations, the stimuli were converted to the HSV colorspace and then rotated on the HSV hue spectrum by a random value between 0° and 360° . Figure 7.1 shows three natural images that have had their hues rotated in this manner. The rotation preserves the saturation of the color in the original input stimuli so that achromatic stimuli remain achromatic, and strongly colored stimuli remain strongly colored. In addition, the number, magnitude, and frequency of transitions between colors are also preserved.

7.2.2 Results

As in previous simulations, color-selective and orientation-selective neurons organize into color-selective and orientation-selective maps in V1 and V2. Most neurons are either color or orientation-selective, and neurons between two types of feature-selective regions are selective for both features. Maps of V2 are not shown in this section since their color-selectivity maps are always a subset of V1 color-selectivity maps, and do not contribute any additional results.

As in Chapters 5 and 6, there are three dominant types of neurons: blue, red, and green-selective. Red and green-selective blobs (shown in Figure 7.2) regularly touch each other and form yellow-selective regions. These yellow-selective regions are slightly more pronounced than they are in the simulations from previous chapters. The histograms of the maps, also in Figure 7.2, reveal that yellow-selective neurons are as numerous as the red or green-selective neurons. However, they do not appear as solitary blobs but are only found next to red and green-selective areas.

Like previous simulations, there are many more blue-selective neurons than any other neuron type. Likewise, these neurons occasionally interact with green-selective blobs to create cyan-selective neurons. Though red and green blobs regularly touch each other to form yellows and oranges, it is much rarer for a blue region to create cyan-selective and magenta-selective areas. Red and green cone sensitivities overlap a great deal, which may be the reason why red and green-selective neurons interact more than the blue-selective neurons. Anecdotal evidence throughout the project also suggests that a coextensive short-cone receptive field causes red and green blobs to be more likely to touch and form yellow-selective regions, but it is not clear under what circumstances this property holds. In one elongated color-blob in the lower-left quadrant of Figure 7.2b, a band of colors is seen: red to yellow to green to cyan to

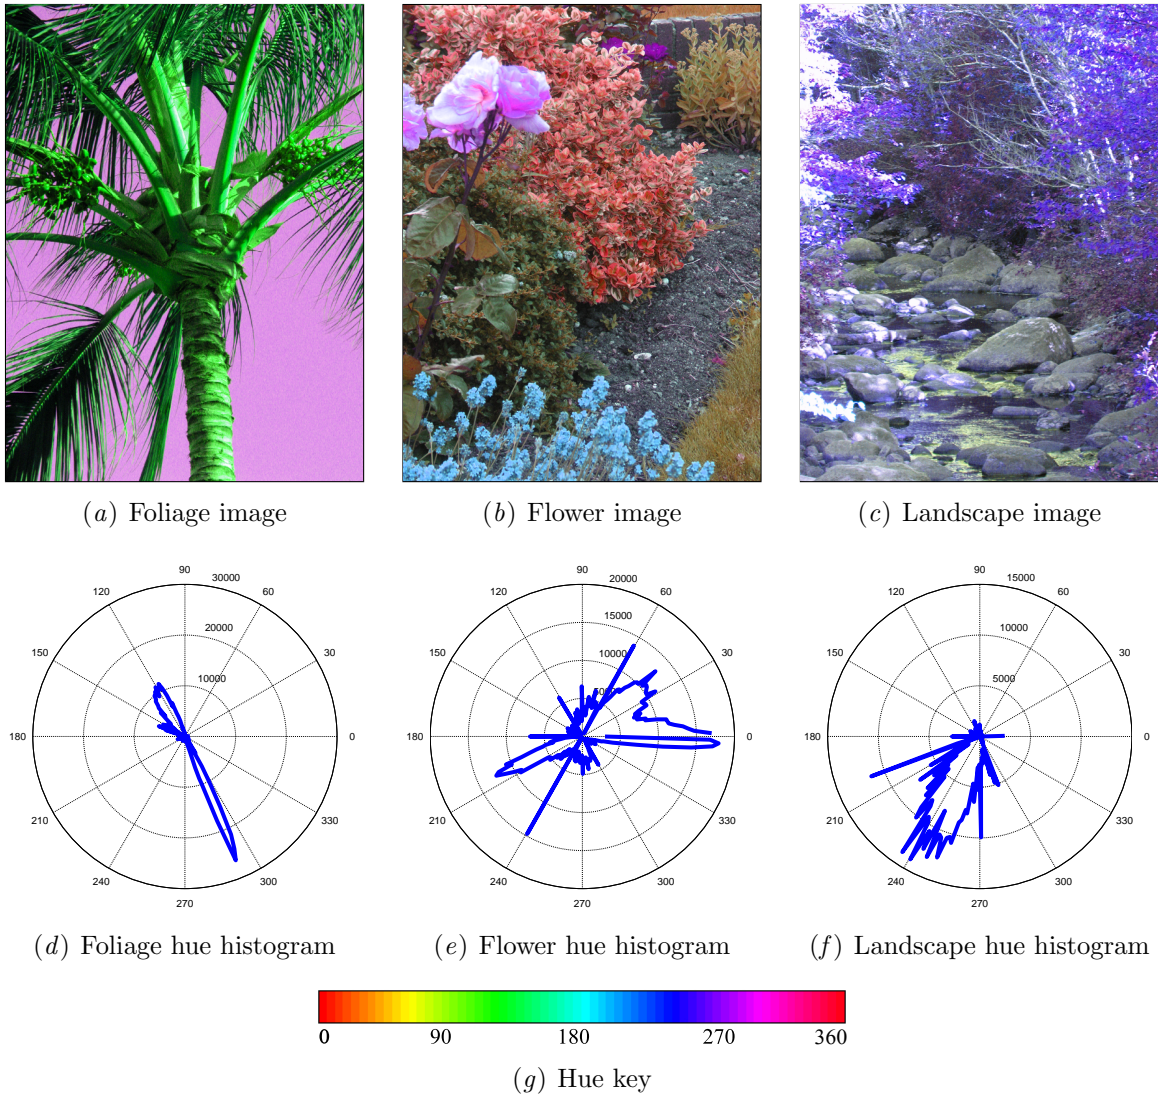


Figure 7.1: **Natural image training stimuli with rotated hues.** Plots (a)–(c) are three sample training images, each from a different image corpus. They have had their hues rotated 220° , 90° , and 300° , respectively. Plots (d)–(f) are polar plots of the colors found in each rotated hue image, and (g) is the color key for the polar histograms, starting with red at 0° moving through all the colors, and then back to red at 360° . For comparison, these three images are also shown with their natural colors and histograms in Figure 4.3 of Section 4.3. Each of the images retain their original hue saturation values, quantity of color transitions, and unbalanced histogram distributions, but a training corpus that contains these images will have randomized and unnatural hue distributions.

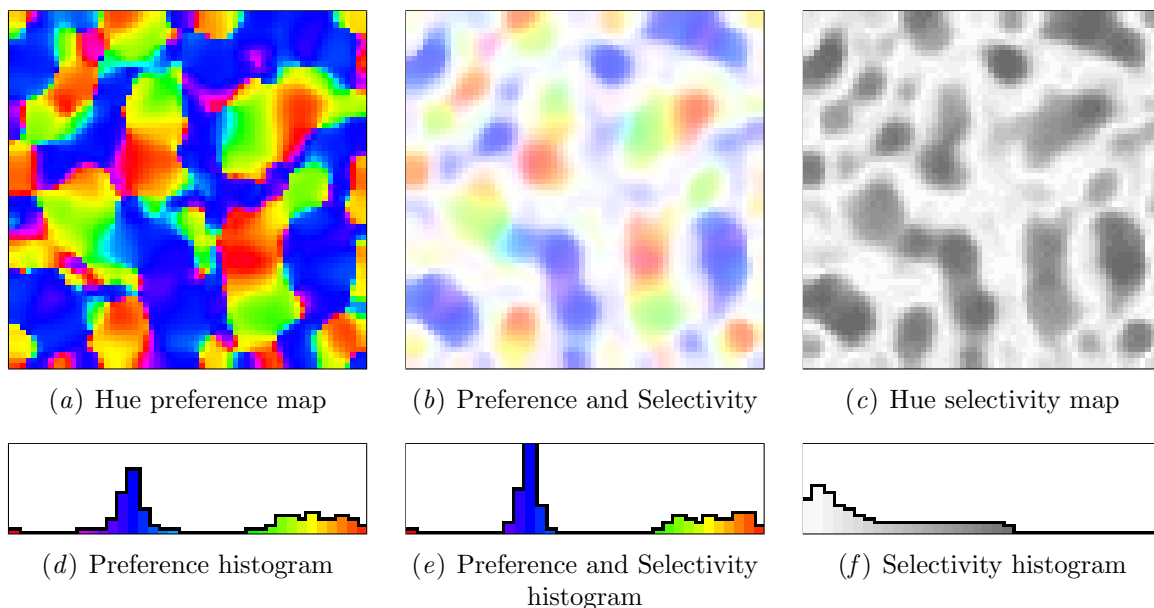


Figure 7.2: **Color-selectivity maps of V1 trained with rotated hues.** V1 maps self-organized on stimuli with random hues develop color-selectivity maps similar to the simulations presented in Chapter 5. Therefore, natural image color biases have limited effect on the hue preferences of V1 maps. Each of the 64×64 neurons are colored with its preferred color in (a) and (b). Red and green blobs regularly interact to form yellow selective neurons. Plot (c) shows the color selectivity of the V1 neurons with regions unselective for color marked as white. The color-unselective regions match regions selective for orientation (not shown). Plot (b) shows each neuron’s preference along with the selectivity, again with color-unselective regions in white. Plots (d)–(f) are histograms of the V1 neurons in (a)–(c), and show that the map is dominated by a large number of blue-selective neurons followed by equal numbers of red, yellow, and green-selective neurons.

blue to purple. Similar to the simulations from previous chapters, magenta selective areas do not appear robustly.

This simulation gives three important results: First, the natural image simulations presented in previous chapters—complete with color biases—have color-selectivity maps that are similar to maps generated with random colors. Second, without some type of photoreceptor channel gain, brightness control, or perhaps just fewer blue cones in the retina, blue-selective neurons will dominate the network. And third, random-hue inputs increase the chance of secondary colors forming, such as yellow, cyan, and purple, bringing the model closer to biologically realistic color maps that contain all possible hues.

Overall the results strongly resemble the Chapter 5 model trained on natural

images. These results suggest that color-preference maps are rarely affected by the strong color biases in natural images, and that on a continuum between color maps with no color-selective neurons and color maps with all possible color preferences, simulations trained on random colors and natural images fall in the middle. Therefore, as will be discussed in Chapter 8, either the natural image corpora in this study do not contain some aspect of the biological color experience, or something must be added to the model if it is to develop full-hue V1 color maps from natural images.

The next section presents the most biologically realistic V1 and V2 maps which have color-preference maps that contain all possible hues.

7.3 Self-organizing the full color spectrum

Simulations regularly create red, green, and blue color blobs, and to a lesser extent secondary colors such as yellow, orange, and cyan. Yet one color is conspicuously absent: magenta (blue plus red). Why is magenta missing from the model when apparently V1 and V2 contain all possible hues (Xiao et al., 2007, 2003). The simulation in this section shows that the LISSOM model is capable of developing color preferences for all hues in the color wheel. Magenta is missing because the simulation is presented with natural images that do not contain enough magenta and color transitions containing magenta.

7.3.1 Training stimuli

As in the previous section, the network was presented with 23 images of natural scenes, where the colors within each image had been randomly rotated. In the last section, the hues in the images were randomly shifted before every presentation to the retina, for a total of 20,000 rotations. In this section, the hues in each image were randomly rotated a single time for the entire simulation, for a total of 23 rotations. By reducing the number of rotations, color biases introduced by the image corpus are shifted but not so much that they disappear completely.

This process changes the distribution of colors presented to the simulation, while simultaneously preserving the number and type of color transitions. Many smooth color transitions before the transformation will result in many smooth transitions after the conversion, but the starting and ending colors of the transition will be different. If some colors, like magenta, rarely appear in the image, or rarely have transition regions, then after the transformation they appear more frequently and share some of the transitions previously dominated by other colors.

7.3.2 Results

Figure 7.3 shows the resulting V1 color-selectivity map, and Figure 7.4 shows the V2 color-preference and color-selectivity maps. The V2 hue-preference map contains the usual features such as pinwheels, linear zones resembling rainbows, and fractures, suggesting it is properly self-organized. The hue-preference histograms for both V1 and V2 have 100% coverage of the hue spectrum, including preferences for magentas and purples. Magenta selectivity is a new result compared to the previous simulations presented, and is in accordance with what Xiao et al. (2007; 2003) reported.


Another new result is that the simulation generates color wheels in the color-preference maps. In at least two locations in Figure 7.4a, all color preferences come together at a single point and form a ring of color, analogous to pinwheels in orientation maps. In further analogy to orientation pinwheels, the color wheel centers are not in color-selective areas. Rather, color-selective blobs tend to have linear color transitions over small pieces of the spectrum.

7.4 Discussion


This chapter first discussed how features in the input stimuli affect the self-organization of feature maps. Color maps form patchy regions of selectivity because color is not as strong an input feature as orientation. However, if the colors are too weak, then no color-selective regions will form at all, or if two color channels have a correlation greater than 99% then the model cannot differentiate between them. At the other extreme, if hue gradients in the images are too strong then color-selective regions in the color maps will become too large to be biologically realistic. Similarly, when colors are too strong, then color blobs will also overlap ocular dominance stripe boundaries, contrary to biology. A balance must be struck between the three types of feature maps for biologically realistic cortex to develop.

Second, simulations trained on artificially adjusted colors reveal that a V1 cortex trained on natural images has similar color-selectivity preferences to cortex trained on random colors. This result suggests that instead of using strictly natural image stimuli in future work, it may be possible to substitute random color stimuli to achieve biologically realistic maps. These simulations also suggest that blue-selective blobs self-organize differently than the red and green-selective blobs. Most interesting, is that as in previous simulations randomly colored stimuli did not form magenta-selective neurons like those found in the brain.

Third, a simulation was presented that developed color-preference maps that contained neurons preferring each of the hues in the color wheel, including magenta. Visual inspection of the training corpus for this simulation showed that there were



(a) Hue preference and selectivity



(b) Hue histogram

Figure 7.3: V1 hue-selectivity map self-organized using artificial-color images. This V1 simulation is unique in that the hue-selectivity map has neurons maximally selective for each of the possible hues. Each of the 64×64 neurons are colored with its preferred hue in plot (a), with regions unselective for color marked in white. These color-unselective regions match regions selective for orientation (not shown). Plot (b) is a histogram of the neuron color preferences in (a), showing that the map contains neurons preferring each of the possible hues. This V1 color-selectivity map matches the biological measurements of Xiao et al. (2007).




Figure 7.4: **V2 color maps trained with artificial-color images.** This V2 region is unique since the hue-selectivity maps have neurons selective for all possible colors. Each of the 80×80 neurons is colored with its preferred hue in plots (a) and (b). Plot (c) shows the color selectivity of the V2 neurons with regions unselective for color marked in white. Plot (b) shows each neuron’s preference along with the selectivity, again with color-unselective regions in white. These color-unselective regions match regions selective for orientation (not shown). Plots (d)–(f) are histograms of the V2 neurons in (a)–(c), showing that the maps contain neurons preferring each of the possible hues. The full-hue color maps are a validation of the model.

many blues, purple, and magentas. The LISSOM model requires biased hues to develop preferences containing the complete color wheel, but the biases must not be so large as to affect the self-organization of other feature maps.

7.5 Conclusion

The results presented in this chapter show that through self-organization the LISSOM model creates cortical maps containing preferences for all the hues in the HSV color space, as well as bands of color selectivity that form gradual transitions between color preferences. Since animal cortex also contains gradual color-preference transitions and broad color preferences (Xiao et al., 2007, 2003), this chapter further validates the accuracy of the LISSOM model. The next chapter will discuss in greater detail the

results of the model, as well as offer some promising directions for future work.

Chapter 8

Discussion and Future Research

The last four chapters presented the results of the LISSOM trichromatic visual cortex model, and shows the details of the self-organized cortex both at the map and the neuron level. This chapter evaluates the implications the results have for the fields of computer science, neuroscience, and psychology. The chapter is divided into six sections. The first section briefly reviews how the dissertation contributes to science in general. The second section evaluates the LISSOM model and discusses future work on it. Third, the natural image stimuli are discussed. The fourth section discusses the lessons learned using natural images as input to a self-organizing system and proposes future work in this area. The fifth and sixth sections evaluate the model V1 and V2 results and propose future extensions and biological experiments.

8.1 Contributions of the dissertation

The dissertation furthers our knowledge of computer science and contributes to both neuroscience and psychology. First, the model provides computer science in general, and robotics in particular, with a biologically inspired color vision architecture that is both robust and scalable. The neural network simulation is a complex synergistic mechanism with interesting mathematical properties in its own right; meanwhile, having this model of color processing enables researchers to develop and study real-world applications that depend on color.

Second, the dissertation contributes to the study of the brain. The many specific predictions presented in previous chapters and also discussed below help direct future experiments and suggest new avenues for scientific exploration. The descriptive power of the model provides a framework to discuss the vision system and how biological mechanisms interact. In this dissertation, the theory of such processes is embedded in a computer simulation. The theory then is not only an intellectual

exercise but also a program that works and duplicates the processing of the visual system.

Third, psychology now has a biological model of color vision to help study perceptual phenomena. For example, optical illusions such as the McCollough effect (McCollough, 1965) or shape from shading illusions (Kingdom, Rangwala, and Hamamji, 2005) require a color model such as this one to test competing hypotheses. The more biologically realistic the model becomes, the better positioned the computational model is for studying additional cortical regions that receive input from the retina.

8.2 The trichromatic LISSOM model

The LISSOM cortex models how retinal stimuli organize the adult visual cortex during development. The model architecture is adaptable to variations in the input and extends well to new feature types. Below, the strengths and weaknesses of the LISSOM model are discussed first, followed by potential directions for future work.

8.2.1 LISSOM architecture

The LISSOM model provides a useful platform on which biological hypotheses can be tested. The input statistics and the visual system organization can be studied in ways not possible in biological experiments or in non-developmental models. Every neuron can be examined, and every connection measured, without disturbing the system.

The model is also a robust system that scales well and allows many overlapping feature maps. However, as the number of feature maps in the cortex increases, there is a systematic and gradual degradation in the feature map quality. There are a fixed number of neurons in the cortex and as more feature maps are added neurons become divided based on preferences for different features in the input. There is no cataclysmic failure of the system, rather the self-organized maps begin to have fewer well-defined structures. As an example, a one-eye achromatic model will develop extremely strong orientation-selectivity maps. If the simulation is extended to two eyes, then ocular dominance maps form but at the cost of the orientation-selectivity map: neurons begin to become partially eye selective along with being orientation selective. As additional features such as color or motion are added, the neurons continue to be selective for even more features and the maps will continue to become less well delineated. Therefore, the model could be extended to increase the quality of feature maps when the retina is presented with multiple features.

8.2.2 Future work

There are a number of possible extensions to the LISSOM model. They can be divided into six categories: (1) increasing the realism of the model, (2) open questions in biology, (3) visual illusions, (4) robotic vision, (5) modeling other animals, and (6) modeling visual regions beyond V1 and V2. Each direction will be addressed separately in this section.

There are two main ways in which the model can be made more realistic. First, the neurons in the modeled cortex are simple cells with well-defined response properties. In the biological cortex there are also neurons that respond in ways that the simple cell model does not, such as activating uniformly across all phases of stimuli in their receptive fields. Including such complex cells into the LISSOM model will make the model more biologically realistic and make it possible to study how complex cells form and interact with simple cells. One possible way that neurons develop complex cell responses is through groups of simple cells linked together in ordered patterns (Chance, Nelson, and Abbott, 1999). Sit and Miikulainen (2007) already extended LISSOM with such a model for monochromatic complex cells, but further work is necessary to extend it to color.

Second, to study map formation more robustly, a 3D cortex needs to be created where the individual lamina of the cortex are implemented. With more than a single layer of cells, each vertical penetration of the cortex measures more than one neuron, and the average activation from all the neurons in the penetration can be used to measure preference maps. Such layers may increase the quality of feature maps since neurons at different layers in the cortex could self-organize selectivity for different features, and allow the cortex to be selective for more features with less feature-map degradation.

There are three directions of future work to help answer open questions about biological trichromatic vision. First, as mentioned in Section 2.2.2, the shape of blue-yellow receptive fields is still controversial. Does the short-cone vision pathway use a coextensive or a difference-of-Gaussian receptive field architecture? Center-surround receptive fields act as a mechanism for automatically balancing luminosity, and may play a role in the activation-balancing mechanism reported by Klug et al. (2003). The models presented in this dissertation are based on the coextensive model which is the most probable given the biological evidence, but additional work should be done to study the differences between the two competing approaches.

Second, the ratio between the long and medium photoreceptors in the retina is known to vary widely between subjects and yet normal color vision is preserved (Brainard, Roorda, Yamauchi, Calderone, Metha, Neitz, Neitz, Williams, and Jacobs, 2000). The mechanisms that compensate for these variations and the process through which they affect cell ratios in the LGN and visual cortex are still not fully understood.




Figure 8.1: **Retinal cone geometry.** How are receptive fields constructed by LGN cells? One hypothesis is that cells randomly group cones in the retina for their RFs. This figure shows a schematic diagram of the retina and how photoreceptors that are spatially near to each other can be combined into receptive fields. The two patches of light-colored photoreceptors indicate how color-opponent receptive fields can be created by grouping together a region of photoreceptors containing different types of cones. Future experiments will include adjusting the ratios of photoreceptor types to examine how the random-wiring hypothesis affect cortical self-organization. Figure adapted from Lennie (2000).

Using LISSOM, it is possible to explore the effects of changing photoreceptor ratios, and perhaps demonstrate a self-organizing mechanism for preserving normal color vision across subjects.

Third, there are two theories about how the receptive fields of LGN neurons form (reviewed by Gegenfurtner and Kiper, 2003). The first is that an ordered mechanism exists to systematically connect photoreceptors into the achromatic and color-opponent receptive fields of LGN cells. The second is that the receptive fields are randomly constructed based upon spatial location. Figure 8.1 shows how receptive fields can be constructed from a random distribution of photoreceptors in the retina by pooling activation from photoreceptors that are spatially near each other. The ordered and random-wiring model can be implemented in LISSOM to make predictions about which model creates maps that best fit the biological data.

The third main direction for future work involves studying visual illusions. Julien Ciroux (2005) modeled the McCollough effect using a simpler dichromatic ar-

chitecture originally developed for the early stages of this dissertation. This previous success of LISSOM suggests that other color-based visual illusions can be modeled using a trichromatic LISSOM model. One visual illusion in particular is a prime candidate to be modeled: An appearance of depth results from viewing a luminance grating and a chromatic grating when they have different orientations (Kingdom, 2003). It would be interesting to see if a linkage between depth and color selective neurons comes out of the self-organization process in LISSOM, and reproduces this depth illusion. However, the model would first have to be extended to include disparity selective inputs.

The fourth category of future work is robotics: Now that the model has been shown to match animal data, it would be interesting to adapt it for use in the early stages of vision in robots. Robotic visual systems generally work differently from the human visual system: Fourier transforms do not process images the same way color-opponent LGN cells respond to photoreceptors. It would be helpful for both systems to process colors in the same way. Moreover, human vision is superior to robotic vision, and by studying the differences between the color vision model and current color processing in robots it may be possible to develop more robust robotic vision algorithms.

The fifth direction is to model other animals. The simulations in this dissertation focused on the macaque monkey as a model for human trichromatic vision, but it is possible to use the model to study how color vision is implemented in various differing species. For instance, dichromatic animals have substantially different functional map structures; such structures should come out through self-organization as well. In each case, a model of a new species can be constructed by changing the photoreceptor activation functions and the LGN architecture; the V1 and V2 architecture would remain largely similar.

In Japanese monkeys proper color perception is dependent upon the environment in which they are raised (Sugita, 2004). It would be fascinating to measure the color-selectivity maps of these monkeys, and model their development computationally. More experiments that study the color-selectivity of developing animals would help our understanding of color map self-organization. In particular, it is not yet clear what prenatal color maps look like, and how these maps change as the animals mature.

The sixth direction is to extend the model to other regions. The modeled color-selective V1 can be used as input to regions beyond V2. In particular, V4 has for a long time been considered a part of the brain's color processing system, and it receives stimuli from a variety of regions including V1 and V2. An additional LISSOM sheet can be added to simulate V4, which would help explain how a vision hierarchy emerges, and how V4 contributes to color processing. All the future work described in this section will help us to further understand how color selectivity develops in the

brain, as well as how the brain integrates and uses color for perception.

8.3 Training stimuli and natural images

Natural images were originally selected for network training since it was not clear what artificially constructed color stimuli should look like. Through the experience gained by training the model with natural images, we now know much more about how natural images affect the computational model. This section will first discuss the lessons learned by using natural images, and then the most promising future work suggested by this research.

8.3.1 Colors in the World

Oriented edges, or luminosity contours, are a nearly universal and strong feature of the natural images. Since natural images contain objects such as foliage, mountains, water, and rocks, regular light variations are expected. However, even natural images do not perfectly mirror the world. The earth is full of areas with slow luminosity transitions, like clear blue sky, or large areas of blank rock. These flat areas fail to make interesting pictures and therefore are left out of image corpora. But the missing blank areas may not be important since it is the distribution of what humans and animals look at that affects self-organization of the visual cortex. Bright colors draw the eye, so perhaps our visual experience has more colors than the world as a whole.

Images with strong luminosity transitions are best for training LISSOM, which requires variations in the input stimuli to drive cortical learning. If images with fewer edges are used for training, then more iterations must be done to complete the self-organization.

The image corpora studied in this dissertation have strong color biases: greens and reds are extremely prominent in the images, with blues, purples, and magentas much less common. Certain hues of blue are often present when sky is seen through the leaves of trees or reflected off the surface of water, and apparently this amount of blue is enough to consistently self-organize three-color selectivity. But these natural image sets do not self-organize magenta color-selectivity in the same manner that the artificially modified stimuli does. Is this because the image corpora do not have enough purples and magentas? Or is there another reason? It will be interesting to try to answer these questions in future work.

8.3.2 Future work

There are three directions for future work with natural images. First, it has not previously been possible to come up with guidelines for constructing colored artificial stimuli that mimic natural images. Training sets that consist of actual natural images can be difficult to work with, since the individual images selected for training can have biases and perturbations that do not mirror the more general natural image statistics. Also, many images contain large blank areas that do not cause network activation. If the simulation presents patches from these blank areas then it takes more presentations to complete the network self-organization. By guaranteeing retinal stimulation, artificial stimuli will make cortical maps self-organize more quickly, saving computer resources and allowing more complicated simulations.

The dissertation suggests a recipe for creating artificial color stimuli: Most importantly, achromatic gradients need to dominate over hue gradients in these stimuli. Simultaneously, there must be enough hues to make each photoreceptor type have a minimum unique entropy (as measured in Chapter 4) of greater than three bits per pixel. Curiously, the actual shades of the colors presented may be relatively unimportant. Colored Gaussians are one possible stimuli that may provide sufficient orientation information while providing fine-grained color control. The saturation and variety of colors in the Gaussians need to be balanced so that the size of color-selective regions are within biologically reasonable sizes. Such a training set would likely generate stronger selectivity maps and allow experimenters to have more control of the training inputs. Better feature maps, and training stimuli that are easier to manipulate, will allow building models of higher cortical regions that use V1 and V2 maps as input, such as V4.

Second, the LISSOM map measurement procedure can be converted from the HSV colorspace into a more biologically realistic colorspace. Two such candidates are XYZ, which has color axes based on human color perception (Vos, 1978), and LUV, which attempts to have colors be perceptually uniform across the color space (Alman, Berns, Snyder, and Larson, 1989). The HSV color representation used in this study is helpful since all hues can be represented as a single real number between 0 and 1. However, the distribution of colors for the hue component of the pixel does not match the distribution or discrimination features of the eye. For example, varieties of yellow in the HSV space are underrepresented compared to reds. It is dangerous to make strong claims about fine color-selectivity variations while map measurements are performed in a color space suited more for the computer than for psychophysical studies. Changing to other color spaces will change color representation from the current one degree of freedom to two or three degrees of freedom. This modification will increase the complexity of color measurements and plots, but it will also increase the biological realism.

The third possible direction of future work is an interdisciplinary project to construct new image corpora that focus on colors. Many of the existing publicly available sets focus on locations (Uganda), artistic quality (Visual Delights), or objects (McGill). The McGill corpus is the best so far in that the variety allows pooling different types of images, but there is no emphasis on capturing the variety of colors in the world, or studying the color statistics in different environments. The color LISSOM project would directly benefit from more and better color corpora, and so would disciplines such as neuroscience and psychology. Neuroscientists need to have precise stimuli to present to experimental animals; psychologists who study human perception need images of natural scenes that humans perceive. It would be an interesting and productive venture to create an interdisciplinary project with either of these fields to further study the natural world and the colors in it.

8.4 Primary visual cortex

It is a validation of the model when features of the simulated V1 match features in biological V1. Similarly, newly observed features of the simulated V1 generate predictions that can be tested in biological experiments. This section presents four predictions of the V1 cortex, then discusses future work.

8.4.1 Predictions of the model

The V1 model generates four predictions about the biological cortex.

Prediction 1: Three types of color selectivity

The main prediction of the model is that there are three major types of color-selective blobs: red, green, and blue. Each of these three types have neurons with a unique arrangement of connections to the LGN regions, so that red-selective neurons are found in red-selective blobs, green-selective neurons in green-selective blobs, and blue-selective neurons are in blue-selective blobs. This prediction means that V1 is able to distinguish among, and is partially selective for, colors that match the long, medium, and short cone photoreceptors in the retina.

Prediction 2: Color-selective neurons for yellow, cyan and magenta

When green-selective and red-selective blobs are adjacent, the neurons between them become selective for yellow, with a gradual transition through the color spectrum from green to red. Similarly, when artificial stimuli are used for V1 training, a large

number of neurons selective for magenta and cyan form in the same way between blue and red, and blue and green blobs. This result predicts that neurons selective for secondary colors form when color blobs are spatially adjacent in biological cortex, and that the full range of hue preferences will develop, based on mixtures of red, green, and blue blobs.

Prediction 3: Color affects eye-preference maps

The amount of color in training images influences where color blobs appear in the ocular dominance maps. Increasing the amount of color in training images increases the likelihood that the cortex develops abnormally: Color blobs are found in binocular regions of the cortical maps and thus cross the left-eye and right-eye selectivity stripes. The reason is that the average color gradient in such training images is greater than the luminosity gradient. The model predicts that abnormal blobs will occur if a developing animal is subjected to stronger colors than found in nature. Similarly, for animals that pay more attention to color cues, the color blobs will not have the same relationship with ocular dominance maps as they do in macaque monkey.

Prediction 4: Lateral connections of color-selective neurons

Neurons selective for color connect to other color-selective neurons with similar chromatic preferences. Red-selective neurons strongly connect to other red-selective neurons, green-selective neurons connect to green-selective neurons, and blue-selective neurons to blue-selective neurons. Such connection patterns constitute a prediction of the model.

8.4.2 Future work

There are five immediate directions for future work. First, the existing V1 model can be modified to explore photoreceptor ratios found at the fovea/periphery boundary. Unlike the periphery, the fovea does not have short-cone photoreceptors, and it is not clear how the missing cones change the cortical organization.

Second, color blindness can also be modeled. LISSOM is a developmental model of color-selectivity, and color selectivity depends upon environmental stimuli in infancy (Sugita, 2004). It may therefore be possible to build a predictive theory about how color-blobs form abnormally by radically modifying the network training stimuli.

Five to ten percent of the population is colorblind to some degree (Sharpe, Stockman, Jagle, and Nathans, 1999). The most common color blindness is caused by a genetic mutation, and results in something similar to dichromatic vision, where

the patient cannot tell the difference between colors that normally activate the long and medium cones. Other subjects have medium and short cone pigments that are too similar to each other and are treated as a single cone type by the brain. Still another type of color blindness is caused by patients having only a single cone type resulting in monochromatic vision. All of these types of color blindness can be implemented with the model to study their effect on cortical self-organization.

Third, the interaction between color and ocular dominance maps can be explored further. As was discussed above, the correlations between the retinal inputs determine where the color blobs develop, but further work needs to be done to determine more specifically what drives the transition between the normal and abnormal regimes of self-organization.

Fourth, the effect of color blob preferences on psychophysical studies can be explored. In previous chapters the color blobs were categorized according to the color of their maximum sensitivity, and consequently were labeled red, green, and blue. Interestingly, these hues do not match to the peak sensitivities of the photoreceptor cells, but are closer to the wavelengths generally considered to appear red, green, and blue to an observer. One interpretation of this observation is that the preferred colors of each blob correspond to the psychophysically defined primary colors of light. Given that humans share the same sensitivity correlations among their the three types of photoreceptors, and that the center-surround receptive fields of the LGN region exist in red-green and blue-yellow combinations, it is likely that the three dominant colors in V1 are a robust feature of biological cortex. That is, the color preferences of each blob type will be the same across different individuals. Future work can establish whether there is indeed a standard color preference for each blob type in the cortex, and if so how it depends on the photoreceptor sensitivity functions and ganglion-cell/LGN processing.

Fifth, it will be interesting to see which color-selective neurons activate when the same object is presented under different illumination. Perception of color remains constant under widely varying illumination conditions, and the model can be examined to see how V1 might play a role in this process. Similarly, LISSOM can now model the perceptual effect of color adaptation, where the cortex becomes less sensitive to colors over time (Engel and Furmanski, 2001). Delahunt and Brainard (2004) report that human color constancy does not map to natural daylight statistics, so it will be useful to see if LISSOM has similar color constancy features. Each of these directions for future work described above would increase the realism of V1, and create a model with even more predictive power.

8.5 Secondary visual cortex

The modeled V2 has topographically organized color-selective and orientation-selective neurons that form specific lateral connections to other V2 neurons. The feature preferences of V2 neurons match biological data, while the lateral connections of the model are predictions. However, the behavior of the model V2 is not as clear as for V1, as described below.

8.5.1 Inputs and color wheels

Curiously, color maps with neurons selective for all hues formed only when the colors in the training images had been artificially adjusted. Color-preference maps containing all possible hues are necessary to develop color wheels, which are important because they are observed in biology (Xiao et al., 2003). Assuming the results from Xiao et al. (2003) are correct, one of two things is happening in the model: Either the natural images used in this study do not have “natural” color statistics matching the long term visual experience of laboratory macaques, or V2 cortex requires additional homeostatic mechanisms that are currently not included in the model. The artificial stimuli of Chapter 7 may be simulating a developmental process that is contributing to full-hue V2 self-organization. Anecdotally, if color-selectivity depends solely on colors presented to the retina then one would expect more regional problems, given the broad color variation in the world (Alaska vs. the Amazon). Color-blindness experiments in developing animals would help settle this question.

This rest of this section will present two predictions of the model before discussing V2 future work.

8.5.2 Predictions of the model

The LISSOM V2 results generate two predictions.

Prediction 1: Color preference of V2 neurons

V2 develops color-selective blobs and is always selective for variations of the color preferences found inside of V1. V2 color blobs tend to be less distinctly red, green, or blue compared to V1, and V2 blobs with different color preferences will often touch and create secondary colors (yellow, cyan, magenta). For example, when a red blob touches a green blob, yellow-selective V2 neurons are found in the middle. The formation of neurons selective for secondary colors through color-blob blending is a prediction of the model.

Prediction 2: Lateral connections of color-selective neurons

Similar to V1 neurons, color-selective neurons in V2 laterally connect to other color-selective neurons. Red-selective neurons connect to other red-selective neurons, green-selective neurons connect to green-selective neurons, and blue-selective neurons connect to blue-selective neurons. These lateral connections constitute a prediction of the LISSOM model.

8.5.3 Future work

Extending V2 will increase the predictive power of the model. First, it is necessary to examine why the model V2 feature maps do not contain stripes. That is, neurons in the model V2 organize color-selectivity and orientation-selectivity maps as expected, but they do not form a striped pattern as in biology. One hypothesis is that V2 requires additional features like motion and disparity as input to the retinas. Without strong orientation stimuli, color-selectivity maps do not form blobs, so it might be that stripes will not form in V2 without some other key input feature. In the future, motion and disparity inputs can be included to see if V2 will develop stripes as a result.

Second, making the model more realistic and biologically detailed will increase the predictive power of the model. For instance, the role of feedback connections going from V2 to V1 can be analyzed this way. These recurrent connections will add memory-like effects which may be necessary for color adaptation and color constancy effects. Sit and Miikkulainen (2006) already extended LISSOM to include recurrent V1/V2 networks using achromatic stimuli, and their model can be extended further to include color inputs.

The LISSOM V2 can also be extended in other ways. Since V2 receives projections from the pulvinar region, the pulvinar needs to be modeled as well so that V2 can have a richer set of inputs. The more biologically realistic V2 becomes, the better positioned the computational model is for studying additional cortical regions such as V4 that receive afferent projections from V2 neurons.

8.6 Conclusion

In conclusion, the V1 and V2 maps self-organize color-selective and orientation-selective neurons with biologically plausible receptive fields and lateral connections. The model performs part of the process that turns undeveloped and homogeneous cortical neurons into structured maps that the rest of the brain can use. The self-organizing model provides computer science with an exploratory framework for com-

puter vision, and provides neuroscience with clear predictions about the structure of biological cortex.

The next chapter will conclude the dissertation by reviewing the purpose and goals of this study along with the major contributions of the model.

Chapter 9

Conclusion

The trichromatic LISSOM model replicates the known data on the organization of color-selective neurons in the cortex, and provides a detailed explanation for how this selectivity can develop through simple learning rules. The model provides concrete and novel predictions for future experiments about what lateral connectivity patterns will be found, and about how the color, orientation, and ocular dominance maps interact.

9.1 Summary

Chapter 2 reviewed the physiology and psychophysics of color vision, as well as the computational models of the visual cortex. The visual pathway starts at the photoreceptors in the retina, goes through the lateral geniculate nucleus, to the primary and secondary visual cortex. The visual system is able to perceive colors that map to single wavelengths of light (e.g. green), as well as colors that exist as products of the visual system, such as magenta.

Chapter 3 presented a model that simulates the early visual system of trichromatic mammals. The modeled retinas contain three types of cells that mimic the sensitivity functions of human photoreceptors. The simulated retina views natural scenes and generates activation values that are processed by a lateral geniculate nucleus region that has topographically organized achromatic and color-opponent receptive fields. The lateral geniculate nucleus cells project into a V1 cortical map, which itself has projections to a V2 cortical map.

Chapter 4 presented statistical analyses of several natural image corpora that were used to train the model. These studies revealed that the three photoreceptor types have extremely high inter-channel relationships. Even with the high correlations, each channel usually has between three and five bits of unique information.

Further, the study found that luminosity contours dominate over color transitions in most natural scenes.

In Chapter 5, the simulated primary visual cortex trained on natural images accurately reproduced the neurons and color-selective structures found in biological cortex. Further, the cortex duplicated the relationship found between color blobs and the orientation and ocular dominance feature maps. Along with duplicating experimental evidence about biology, the model also makes a number of predictions about the selectivity of color blobs and how they interact with other neurons.

In Chapter 6, the simulated secondary visual cortex self-organized much like the primary visual cortex. However, V2 also has distinct differences, in that larger receptive fields allow neurons to integrate over a larger region of the visual field. The model also makes concrete predictions about the color selectivity and lateral connectivity of V2 neurons.

Chapter 7 presented further systematic experiments where the colors of natural images were manually modified. If any channel within a training corpus had less than three bits of information per pixel, or if color contours dominated orientation contours, then the map did not self-organize properly. These simulations also revealed that V1 and V2 cortex are capable of developing color preferences for all possible hues in the HSV color-space. These results confirm that the results of the last two chapters are an outcome of the training images and not solely a product of the network architecture.

Chapter 8 discussed the results of the model, and also listed the predictions that the model generated. A computational model of color vision enables future research in robotic color vision, perceptual modeling, and more advanced biological simulations. Further, the most interesting directions for future work were presented which including increasing biological realism, and modeling color blindness.

In summary, the model of the visual cortex created biologically realistic feature-selective maps. In addition, the interaction among the three map types in V1, and the interaction between the two map types in V2, also matched known physiology. The different organizational patterns found in the feature maps are a result of orientation inputs being a stronger feature in the stimuli than color, and the strong correlation of the inputs to the left and right eyes.

Each color-selective blob primarily contains a single type of color-selective neuron, either red, green or blue. The model also suggests how neurons can become selective for other hues such as yellows, magentas, and blues. That is, neurons in areas where two blobs meet respond equally to both blob types and as a result become selective for additional colors. A similar mechanism may take place in biology, giving rise to cortical neurons with a variety of color preferences. Therefore, the color LIS-SOM model helps explain how color vision arises in the brain, and why the cortical maps organize into different patterns.

9.2 Conclusion

Thus the model is an embodiment of our current knowledge about the early visual system and the self-organization mechanisms that creates color selectivity in the cortex. This work is a platform for testing our scientific knowledge and generating predictions to guide experiments. In this way, the model helps us better understand visual perception in general, and the physiological basis for color perception in particular.

Appendix A

Converting RGB images to LMS cone activations

The RGB (red/green/blue) to LMS (long/medium/short) transformation described in this appendix is used to simulate how human L, M, and S photoreceptors will respond to RGB images displayed on a computer monitor. For biological realism, the LISSOM model requires LMS values as input, yet most training data is available only as RGB images. The RGB→LMS transformation allows LISSOM to use widely available uncalibrated color bitmap images as training inputs.

The transformation is necessary because RGB triples are only a coarse approximation to what each of the photoreceptor types will see. For instance, an RGB image containing nonzero values only in the red channel will activate both long and medium-wavelength photoreceptors when displayed on a computer screen, because of the overlap of the sensitivity functions for these two cone types (see Figure 2.3). The following sections go through each of the steps in this transformation, and also shows how computer code to do this transform can be obtained.

A.1 RGB to LMS algorithm

To convert from an RGB image where each triple represents phosphor luminance, to an LMS image where each triple represents cone stimulation, requires three data sets:

1. The wavelength sensitivity function for each type of retina cone.
2. Phosphor photon emission functions for a specific computer monitor.
3. RGB images of natural scenes to be converted.

The emission spectra of a specific CRT monitor must be measured because each monitor has different energy emission functions. Even so, variations between monitors are small enough that any display with realistic looking colors should be sufficient to use in the conversion. The calculated LMS values simply represent one possible high-fidelity rendering of the RGB data by a computer monitor.

The major steps that will be involved in the transformation are: (1) calculating the photon emissions for each monitor gun at every possible pixel intensity, (2) summing the photon emissions for each monitor gun at the specified RGB pixel intensity, and (3) calculating the final cone activity as the dot product of the cone sensitivity function with the summed phosphor emission values.

A single monitor pixel $(\alpha_R, \alpha_G, \alpha_B)$ will be converted to $(\alpha_L, \alpha_M, \alpha_S)$, where α_R is the intensity in the red channel of that pixel, α_G is the intensity in the green channel, and α_B is the intensity in the blue channel. Similarly, α_L is the activation of the long cone, α_M is the activation of the medium cone, and α_S is the activation of the short cone. The same steps will be repeated for each pixel within an RGB training image.

Notation

- All uppercase letters (L, M, S, \dots) represent single dimensional numerical row-major column vectors. Subscripts also denote separate vectors: P_1 is different from P_2 .
- A^T means the transpose of A .
- Lowercase letters (α, i, j, \dots) are scalars or variables.
- All normalizations are ∞ -norms. The ∞ -norm of vector A is the largest element of $abs(A)$.

Step 1: Calculate the monitor spectrum emission functions

It is necessary to normalize the maximum spectral power distribution vectors to use the full dynamic range of the pixel storage formats. When normalizing, the relationship between emission functions must be preserved because the relative energy levels affect the photoreceptors and the peak energy of each gun will rarely be the same. Figure 2.3 shows that the red gun has a higher energy peak than the others. Dividing all three vectors by the wavelength with largest energy will create a dataset that has a maximum single-wavelength energy of 1.

The normalizing factor p is

$$p = \left\| \left[\|E_R^{max}\|, \|E_G^{max}\|, \|E_B^{max}\| \right] \right\|, \quad (\text{A.1})$$

and the normalized spectral power is

$$E_R^{norm} = \frac{E_R^{max}}{p}, \quad E_G^{norm} = \frac{E_G^{max}}{p}, \quad E_B^{norm} = \frac{E_B^{max}}{p}, \quad (\text{A.2})$$

where E_R^{norm} , E_G^{norm} , and E_B^{norm} are the normalized spectral power distributions for phosphors at maximum luminosity. The area under each curve will be different, and the maximum value for each curve will be different, as they should be.

Step 2: Total spectral power emission for each image pixel

To calculate the luminance power for a particular pixel, the summation of the light coming from each red, green, and blue phosphor is summed together. CRT monitors emit photons as a linear function relative to the input current, but cone luminosity sensitivity is non-linear. Monitors include an exponential function, usually called the gamma function, to adjust the display. The exponentiation of the pixel values by the gamma function creates a linear increase in perceptual luminosity as the RGB values increase.

The exact gamma function may vary depending upon the monitor, but a standard equation is: $\gamma(x) = x^{2.2}$. The top plot of Figure 2.3 on page 10 shows the emission vectors for when red, green, and blue are maximum ($\alpha_R = 1, \alpha_G = 1, \alpha_B = 1$). By using the gamma function $\gamma(x) = x^{2.2}$; $x \in [0, 1]$ and the maximum spectral power distribution vectors E_R^{max} , E_G^{max} , and E_B^{max} (in Figure 2.3), the emission vector for the RGB pixel ($\alpha_R, \alpha_G, \alpha_B$) is calculated as

$$P = \gamma(\alpha_R)E_R^{norm} + \gamma(\alpha_G)E_G^{norm} + \gamma(\alpha_B)E_B^{norm}. \quad (\text{A.3})$$

Each of the three terms $\gamma(\alpha_R)E_R^{norm}$, $\gamma(\alpha_G)E_G^{norm}$, and $\gamma(\alpha_B)E_B^{norm}$ are vectors containing the strength of the phosphor emission for discrete wavelengths at the specified phosphor intensity.

Step 3: Cone activity in the retina

The raw photoreceptor wavelength sensitivity functions L , M , and S must be normalized so that the final activation values will be in ranges usable by the LISSOM model and be able to take on the maximum possible range of values. The normalization method used here assumes that each cone type has an independent gain control

so that when looking at a white light, each cone type becomes fully driven. Each cone sensitivity function is independently normalized so that white light from the computer monitor causes an activation value of 1. If the power spectrum of a pixel showing white light is denoted as P_{max} , then the formula for the normalization is

$$L_{norm} = \frac{L}{L^T \cdot P_{max}}; \quad M_{norm} = \frac{M}{M^T \cdot P_{max}}; \quad S_{norm} = \frac{S}{S^T \cdot P_{max}}. \quad (\text{A.4})$$

The L_{norm} , M_{norm} and S_{norm} vectors replace L , M , and S in future calculations since the normalization allows each cone type to take all possible values, and a white light ($\alpha_R = 1$, $\alpha_G = 1$, $\alpha_B = 1$) generates the maximum cone activation ($\alpha_L = 1$, $\alpha_M = 1$, $\alpha_S = 1$.)

Measuring the cone activity in the retina is done by taking the dot product of the spectral sensitivity vector of the cone (L_{norm} , M_{norm} , or S_{norm} in Figure 2.3), with the emission power spectrum P of the pixel in the RGB image

$$\alpha_L = P^T \cdot L_{norm}; \quad \alpha_M = P^T \cdot M_{norm}; \quad \alpha_S = P^T \cdot S_{norm}. \quad (\text{A.5})$$

Then α_L , α_M , and α_S are the scalar long, medium, and short cone activation values for the image pixel.

A.2 Python conversion code

Python code that performs the above transformation is freely available on the web from the University of Texas at Austin Neural Networks Group:
<http://nn.cs.utexas.edu/keyword?rgbtolms> .

Appendix B

Simulation parameter values

Most of the variables in the LISSOM simulation have standard values derived from a handful of free parameters, and were borrowed from the default monochromatic simulation settings with no modification (Miikkulainen et al., 2005). However, some free parameters within LISSOM had to be adjusted according to the inputs and architecture for each specific model. In addition, five new parameters were added to the LISSOM simulator to support the modeling of trichromatic photoreceptors: Three parameters modify the retinal stimuli, and two parameters specify the start and ending iterations for V2 training.

B.1 Overview

Three distinct simulations were presented in this dissertation. The results of the first were presented in Chapter 5 and Chapter 6 and showed the self-organized feature maps and neurons in V1 and V2. The second simulation studied unbiased training images and was presented in Section 7.2. And the third simulation created color maps with all measured hue preferences and was presented in Section 7.3. LISSOM is robust against variations in parameter settings, and results similar to those presented in the dissertation can be obtained within a range of values.

In this chapter, the parameter values for these three simulations are reviewed. The majority of parameter values are the same for all three simulations, but some parameters have different values between simulations. The first section presents the constant parameters of the color V1 map, which do not change between simulations. The second section presents the new parameters that were added to the retina and V2 regions to support color vision. In the third section, the modification schedules for both V1 and V2 parameters are shown for those parameters that change their value during the training of the model. The fourth section presents the parameter values

Parameter	Value	Description
N_{do}	64	Reference value of N_d , the cortical density
L_{do}	24	Reference value of L_d , the LGN density
R_{do}	24	Reference value of R_d , the retinal density
r_{Ao}	$\frac{N_{do}}{4} + 0.5$	Reference value of r_A , the maximum radius of the afferent connections
r_{Eo}	$\frac{N_{do}}{10}$	Reference value of r_E , the maximum radius of the lateral excitatory connections
r_{Io}	N_{do}	Reference value of r_I , the maximum radius of the lateral inhibitory connections
t_{fo}	20,000	Reference value of t_f , the number of training iterations
w_{do}	0.00005	Reference value of w_d , the lateral inhibitory connection death threshold

Table B.1: **Parameters for the LISSOM simulation.** These values define the base simulation that serves as a basis for other parameters, as specified in Table B.2. The subscript “o” in each name stands for “original”. Adapted from Miikkulainen et al. (2005).

unique to the V2 region, and the fifth section discusses the parameters for both V1 and V2 that change when going from a monochrome to a color simulation.

B.2 V1 and V2 default parameters

Table B.1 shows the reference constants for the derived V1 and V2 parameters shown in Table B.2. The parameter values in Table B.2 are the same for all V1 and V2 simulations presented in this dissertation, except for the modified V2 parameters described in Section B.5. The V1 and V2 cortical regions are separate sheets of LISSOM neurons and therefore each have their own set of parameters within the same simulation.

Previous default monochromatic simulations use two LGN sheets (Miikkulainen et al., 2005), but with trichromatic simulations, the LGN region is expanded to contain 16 sheets, and the memory requirements of V1 neurons increase eight-fold. The greater memory requirements of each V1 neuron necessarily limit the maximum value the V1 N_{do} can take which controls the number of neurons within V1. Similarly, increasing the lateral inhibitory connection radius r_{Io} also increases memory requirements of the V1 cortex, but the increase in radius is necessary to study the long-range correlations and connectivity of color blobs. The other values in Table B.1 are the same as the default monochromatic LISSOM simulation (Miikkulainen et al., 2005).

The majority of the parameter values listed in Table B.2 are also inherited from

the default monochromatic LISSOM parameter values (Miikkulainen et al., 2005). A few free parameters in V1 and V2 had to be adjusted so that the simulation would properly self-organize when trained with color natural images. In particular, the parameters modified were the lower threshold β^i , and the neuron learning rates α_A^i , α_E^i , and α_I^i . Color images of natural scenes create little LGN activation compared to artificial stimuli, therefore the β^i parameter needed to be lowered so that V1 and V2 neurons would be more likely to respond to a retinal stimulus. Meanwhile, the afferent and lateral learning rates were empirically determined so that feature selectivity regions formed smooth transitions following the parameter schedule described later.

B.3 New retina and V2 parameters

To implement trichromatic color vision five new parameters were added to the LISSOM model. Table B.3 shows these new parameters, along with the values they have for the different simulations presented in the dissertation.

The parameters γ_L , γ_M , and γ_S are used to change the brightness of the long, medium, and short photoreceptor channels in the retina. These parameters were introduced because if one type of photoreceptor had a larger average brightness than the other two photoreceptors then the cortex will be biased toward the brighter cone type. Values greater than 1.0 make the training stimuli brighter than the original input bitmap and is usually required for simulated photoreceptor inputs, because the image transformation in Appendix A reduces the average brightness of the images. By adjusting the scaling factor of the three cone types, balanced color-selectivity maps will form in V1 and V2.

The parameter t_{S2} specifies at which iteration V2 will begin training, and t_{f2} specifies how many iterations V2 will train. For example, the values of the simulation parameters for Chapters 5 and 6, shown in Table B.3, will start V2 training at iteration 20,000 and continue for 20,000 iterations. In order for V2 maps to not develop ocular dominance maps, V2 must be trained on monocular stimuli after V1 has self-organized. If ocular dominance maps are not being modeled, then the V2 map can be trained at the same time as the V1 map.

B.4 V1 and V2 parameter schedules

Some V1 and V2 parameters (Tables B.2 and B.3) are constant and do not change during training. Other parameters do change according to a fixed schedule (Table B.4). The parameter adjustments allow early training iterations to form large and broad selectivity blobs in the feature maps, and then have later iterations form fine details within the maps.

Parameter	Value	Description
N_d	N_{do}	Cortical density, i.e. width and height of a unit area of cortex
L_d	L_{do}	LGN density, i.e. width and height of a unit area of the LGN (the area that projects to N_d)
R_d	R_{do}	Retinal density, i.e. width and height of a unit area of retina (the area that projects to L_d)
s_g	1.0	Global size scale of the model in area units N_d , L_d , and R_d
n_A	16	Number of afferent RFs per cortical unit (e.g. 1 ON and 1 OFF)
r_A	$\frac{L_d}{4} + 0.5$	Maximum radius of the cortical afferent connections
r_E^i	$\frac{N_d}{10}$	Initial value for r_E , the maximum radius of the lateral excitatory connections, before shrinking
r_{Ef}	$\max(2.5, \frac{N_d}{44})$	Minimum final value of the r_E after shrinking
r_I	N	Maximum radius of the lateral inhibitory connections
s_w	$\frac{r_A}{r_{Ao}}$	Scale of r_A relative to the default
σ_A	$\frac{r_A}{1.3}$	Radius of the initial Gaussian-shaped afferent connections
σ_E	$0.78r_E^i$	Radius of the initial Gaussian lateral excitatory connections
σ_I	$2.08r_I$	Radius of the initial Gaussian lateral inhibitory connections
σ_c	$0.5s_w \frac{R_d}{L_d}$	Radius of LGN DoG center Gaussian
σ_s	$4\sigma_c$	Radius of LGN DoG surround Gaussian
r_L	$4.7\sigma_s$	Maximum radius of the LGN afferent connections
s_b	1.0	Brightness scale of the retina (contrast of fully bright stimulus)
o_b	0.5	Brightness value of the background of the retina
R_p	L	Width & height of the random scatter of discrete pattern centers
s_s	0.0	Scale of the input pattern scatter from the calculated value
s_d	1.0	Input density scale (ratio between average cortical activity for one oriented Gaussian to the average for the actual pattern)
s_t	$\frac{1}{s_d}$	Iteration scaling factor; can be adjusted to use fewer iterations if input patterns are more dense at each iteration, or vice versa
n_p	$\max(1, s_d s_r)$	Number of discrete input patterns per iteration (e.g. Gaussians)

(Table continues on the next page)

(Table continued from the previous page)

Parameter	Value	Description
N	$s_g N_d$	Width and height of the cortex, in number of units
L	$s_g L_d + 2(r_A - 0.5)$	Width and height of the LGN, in number of units
R	$s_g R_d + 2\frac{R_d}{L_d}(r_A - 0.5) + 2(r_L - 0.5)$	Width and height of the retina, in number of units
s_r	$\left(\frac{L}{L_d + 2(r_A - 0.5)}\right)^2$	LGN area scale relative to the reference simulation
γ_A	1.0	Scaling factor for the afferent weights
γ_E^i	0.450	Scaling factor for the lateral excitatory weights
γ_I^i	0.875	Scaling factor for the lateral inhibitory weights
γ_L	$\frac{2.33}{s_b}$	Scaling factor for LGN's afferent weights
t_s^i	9	Initial value for t_s , the number of settling iterations
δ^i	0.076	Initial value for δ , the lower threshold of the sigmoid activation function
β^i	$\delta^i + 0.55$	Initial value for β , the upper threshold of the sigmoid activation function
t_f	$t_{f_0} s_t$	Number of training iterations
α_A^i	2.78×10^{-5}	Initial value for α_A , the afferent learning rate
α_E^i	4.64×10^{-3}	Initial value for α_E , the lateral excitatory learning rate
α_I^i	2.50×10^{-4}	Lateral inhibitory learning rate
w_d	$2w_{d_0} \frac{r_{I_0}^2}{r_I^2}$	Lateral inhibitory connection death threshold
t_d	t_f	Iteration at which inhibitory connections are first pruned

Table B.2: **Defaults for V1 and V2 constant parameters.** This table specifies how the default values for the V1 and V2 simulations are constructed, based on the reference values from Table B.1. These parameters have constant values in each simulation and are always the same across each of the different simulations. Those with the superscript “i” represent the initial values for parameters without the superscript, that starts at the initial value and then is modified over the simulation, as shown in Table B.4. The table is organized into sections including user-defined network size, connection radius, and input pattern parameters on the previous page, and calculated network size, connection strength, activation, and learning parameters on this page. Adapted from Miikkulainen et al. (2005).

Parameter	Description
γ_L	Scaling factor for long cone gain control
γ_M	Scaling factor for medium cone gain control
γ_S	Scaling factor for short cone gain control
t_{S2}	First iteration of V2 training
t_{f2}	Number of V2 training iterations

Parameter	Chapter 5 and 6	Chapter 7.2	Chapter 7.3
γ_L	1.19	1.40	1.40
γ_M	1.40	1.40	1.40
γ_S	0.70	1.40	1.15
t_{S2}	20,000	0	0
t_{f2}	20,000	20,000	20,000

Table B.3: **New retina and V2 parameters for color simulations.** A color-selective LISSOM model requires adding five new parameters to the simulation. These parameters have different values for each of the simulations presented in the dissertation.

The α and γ parameters within Table B.4 are described in Section 3.1.3. Parameters β and δ are the upper and lower threshold values for the piecewise sigmoid activation function $\sigma(\cdot)$, used in Equation 3.2. The variable t is the number of settling rounds of Equation 3.2 at each time-step. The parameter r_E is the radii of the excitatory connections of V1 neurons as measured in neurons.

The model can properly self-organize maps under a wide range of parameter settings. For example, similar network organizations can be obtained with more or fewer than 20,000 steps, if the parameter multipliers are adjusted accordingly.

B.5 Unique V2 simulation parameter values

In all the simulations, V2 regions have the same values as the V1 regions, except for the parameters listed in Table B.5. The differences between V2 and V1 are discussed here, and are organized into five different categories: cortex size, afferent connection radius, connection strengths, activation thresholds, and learning rates.

First, the number of neurons in the modeled V2 can be larger than V1, because V2 requires much less memory. Therefore s_g was increased to 2.0 which doubles the size scale of the V2 region and makes N become 80. Parameter N sets the number of neurons in V1 to N^2 , and can be increased or decreased based upon CPU memory constraints. The larger number of neurons increases the smoothness of the V2 maps

V1 parameter schedule									
Iteration	γ_I	γ_E	α_I	α_E	α_A	β	δ	t	r_E
0	γ_I^i	γ_E^i	α_I^i	$5\alpha_E^i$	$7\alpha_A^i$	β^i	δ^i	t^i	r_E^i
200	γ_I^i	γ_E^i	α_I^i	$5\alpha_E^i$	$7\alpha_A^i$	$\beta^i + 0.01$	$\delta^i + 0.01$	t^i	$0.600r_E^i$
500	γ_I^i	γ_E^i	α_I^i	$7\alpha_E^i$	$5\alpha_A^i$	$\beta^i + 0.02$	$\delta^i + 0.02$	t^i	$0.420r_E^i$
1000	$2\gamma_I^i$	$2\gamma_E^i$	$2\alpha_I^i$	$10\alpha_E^i$	$5\alpha_A^i$	$\beta^i + 0.03$	$\delta^i + 0.05$	t^i	$0.336r_E^i$
2000	$3\gamma_I^i$	$2\gamma_E^i$	$3\alpha_I^i$	$15\alpha_E^i$	$4\alpha_A^i$	$\beta^i + 0.05$	$\delta^i + 0.08$	$t^i + 1$	$0.269r_E^i$
3000	$3\gamma_I^i$	$2\gamma_E^i$	$3\alpha_I^i$	$15\alpha_E^i$	$4\alpha_A^i$	$\beta^i + 0.08$	$\delta^i + 0.10$	$t^i + 1$	1.5
4000	$3\gamma_I^i$	$2\gamma_E^i$	$3\alpha_I^i$	$15\alpha_E^i$	$3\alpha_A^i$	$\beta^i + 0.11$	$\delta^i + 0.10$	$t^i + 1$	1.5
5000	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	$25\alpha_E^i$	$3\alpha_A^i$	$\beta^i + 0.14$	$\delta^i + 0.11$	$t^i + 2$	1.5
6500	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	$25\alpha_E^i$	$3\alpha_A^i$	$\beta^i + 0.17$	$\delta^i + 0.12$	$t^i + 3$	1.5
8000	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	$25\alpha_E^i$	$3\alpha_A^i$	$\beta^i + 0.20$	$\delta^i + 0.13$	$t^i + 4$	1.5
20000	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	$25\alpha_E^i$	$1.5\alpha_A^i$	$\beta^i + 0.23$	$\delta^i + 0.14$	$t^i + 4$	1.5

V2 parameter schedule									
Iteration	γ_I	γ_E	α_I	α_E	α_A	β	δ	t	r_E
$t_{S2} + 0$	γ_I^i	γ_E^i	α_I^i	$2\alpha_E^i$	$9.3\alpha_A^i$	β^i	δ^i	t^i	r_E^i
$t_{S2} + 1500$	γ_I^i	γ_E^i	α_I^i	$2\alpha_E^i$	$9.3\alpha_A^i$	$\beta^i + 0.06$	$\delta^i + 0.06$	$t^i + 1$	$0.600r_E^i$
$t_{S2} + 2500$	γ_I^i	γ_E^i	$2\alpha_I^i$	α_E^i	$7.0\alpha_A^i$	$\beta^i + 0.10$	$\delta^i + 0.08$	$t^i + 1$	$0.600r_E^i$
$t_{S2} + 5000$	$2\gamma_I^i$	$2\gamma_E^i$	$2\alpha_I^i$	α_E^i	$5.6\alpha_A^i$	$\beta^i + 0.16$	$\delta^i + 0.11$	$t^i + 2$	1.5
$t_{S2} + 6500$	$2\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	α_E^i	$3.3\alpha_A^i$	$\beta^i + 0.21$	$\delta^i + 0.23$	$t^i + 3$	1.5
$t_{S2} + 8000$	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	α_E^i	$2\alpha_A^i$	$\beta^i + 0.33$	$\delta^i + 0.26$	$t^i + 4$	1.5
$t_{S2} + 10000$	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	α_E^i	$2\alpha_A^i$	$\beta^i + 0.36$	$\delta^i + 0.26$	$t^i + 4$	1.5
$t_{S2} + 13000$	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	α_E^i	$2\alpha_A^i$	$\beta^i + 0.37$	$\delta^i + 0.26$	$t^i + 4$	1.5
$t_{S2} + 16000$	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	α_E^i	$2\alpha_A^i$	$\beta^i + 0.37$	$\delta^i + 0.27$	$t^i + 4$	1.5
$t_{S2} + 20000$	$4\gamma_I^i$	$2\gamma_E^i$	$5\alpha_I^i$	α_E^i	$1\alpha_A^i$	$\beta^i + 0.38$	$\delta^i + 0.28$	$t^i + 4$	1.5

Table B.4: **V1 and V2 parameter modification schedule.** The table presents the parameters that change during the network training, and their values at different iterations. The top half shows the parameter values for the V1 region and the bottom half shows the parameter values for V2. Each column shows the value of a parameter at successive iterations during training. The parameter schedule allows the cortex to initially make large adjustments to the network, and then gradually settle toward fine grained organization. The model can properly self-organize maps under a wide range of parameter settings.

Parameter	Chapter 5 and 6	Chapter 7.2	Chapter 7.3
N	80	80	80
n_A	1	1	1
r_A	12.5	12.5	12.5
γ_A	3.0	2.0	2.0
γ_I^i	0.6	0.6	0.6
γ_E^i	1.05	0.55	0.55
α_I^i	5.0×10^{-5}	5.0×10^{-5}	5.0×10^{-5}
α_E^i	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}
α_A^i	1.5×10^{-3}	1.5×10^{-3}	1.5×10^{-3}
β^i	0.726	0.726	0.726
δ^i	0.040	0.040	0.040
s_g	2.0	2.0	2.0

Table B.5: **V2 region parameter initialization.** The values of the parameters in Equations 3.1, 3.2, and 3.3 are changed during training according to a preset schedule. V2 parameter values are the same as the V1 values used in Table B.2 except for a few modified values which are shown in this table. Each cortical region has an independent set of parameters so V1 and V2 are able to have different values for the same parameters within the same simulation.

but does not change the size of the biological cortex that the model represents.

Second, V2 neurons have a retinal receptive field double that of V1 neurons. Given the receptive field size of LGN cells, and the afferent connection fields of V1 neurons, setting r_A to 12.5 gives V2 neurons double the retinal visual field of V1 neurons.

Third, the input connection scale of V2 neurons had to be empirically found, and were set to the values shown in Table B.5. The model self-organizes properly with a broad range of values so these input connection scales can be considered recommended settings for future experiments and are not necessarily required values. Increasing γ_I^i will cause blob edges to form faster, and at large values cause unrealistic fractures to form. Meanwhile, increasing γ_E^i will increase the number and size of activity blobs in the cortical region, and generally must be balanced against the strength of γ_I^i . Also, γ_A affects the strength of the afferent stimuli, and should be increased if the input sheet is not generating enough cortical activity, and decreased if there is too much cortical stimulation.

Fourth, the β^i and δ^i threshold values of the sigmoid activation function needed to be adjusted. All V1 activity is important for training V2 so the δ^i for V2 can usually have a lower value than for V1. Meanwhile, β^i needs to be balanced with the input connection scales so that the normal range of training stimuli causes smooth activity

blobs with gradual changes in neural activity along the blob edges.

Fifth, the learning rates of the afferent and lateral connections have to be adjusted so that learning is not so fast that neurons forget previous presentations, but not so slow as to never become feature selective. The learning rate parameters depend upon how often activity blobs appear in the cortex. All regions of the cortex need to get one or two activity blobs to set initial preferences, then lower learning rates can be used to allow the model to create smooth transitions between the different preference regions.

B.6 Extending a monochrome LISSOM model to color

This section describes further why certain parameters were changed from the values already established as appropriate for the default monochrome simulation reviewed in Miikkulainen et al. (2005).

Cortex size N_{do}

Memory constraints required setting N_{do} to 64. V1 then had 64×64 neurons, each with 16 circular afferent connection fields. Memory requirements increase proportionally to the square of N_{do} . The memory requirements for V2 are less than for V1 since each V2 neuron has a single afferent projection.

Cortical cell inhibitory radius r_{Io}

The default radius of monochrome simulations was not large enough to allow color blobs with the same preferences to connect together. The lateral radius of inhibitory connections was increased to equal N_{do} so that the lateral interactions between color-selective neurons could be established.

Number of afferent receptive fields n_A

Default monochrome simulations have a value of 2, one for the ON LGN sheet, and the other for the OFF LGN sheet. With trichromatic simulations this value was increased to 8, or one each of the ON and OFF versions of the four types of color-opponent receptive fields. With a two-eye simulation, the parameter was doubled to 16.

Inhibitory connection gain γ_I^i

The V1 inhibitory connection gain was reduced slightly to 0.076 and then adjusted throughout the simulation (shown in Table B.4) to keep the activity blobs from developing fractures and sharp edges.

Lower threshold for sigmoid activation δ^i

The lower threshold of the piecewise-linear sigmoid function was reduced slightly from the default monochrome simulation setting because the training images were generally darker than in the monochrome simulations and caused less cortical stimulation. A lower δ^i caused neurons to respond stronger to weaker stimuli. The starting values for δ are shown in Tables B.2 and B.5, while the modification schedule is in Table B.4.

Learning rates $\alpha_E^i, \alpha_i^d, \alpha_A^i$

Color natural images tend to create fewer activity bubbles than monochrome images. To compensate, the learning rates were increased so that the model would learn more from each presentation. Care was taken to ensure that the maps still developed smoothly; increasing these rates too high would result in individual patches learning too quickly, without a global, smooth ordering. Note that the activity patterns affect all of the connection types equally. As a result, the learning rates are tightly coupled and should be adjusted together as a group. The starting values for these parameters are shown in Tables B.2 and B.5, while the modification schedule is in Table B.4.

Bibliography

- Alexander, D. M., Bourke, P. D., Sheridan, P., Konstandatos, O., and Wright, J. J. (2004). Intrinsic connections in tree shrew V1 imply a global to local mapping. *Vision Research*, 44:857–876.
- Alman, D. H., Berns, R. S., Snyder, G. D., and Larson, W. A. (1989). Performance testing of color difference metrics using a color-tolerance dataset. *Color Research and application*, 21:174–188.
- Anderson, J. A., and Rosenfeld, E., editors (1988). *Neurocomputing: Foundations of Research*. Cambridge, MA: MIT Press.
- Barrow, H. G., and Bray, A. J. (1993). An adaptive neural model of early visual processing. In Eeckman, F., and Bower, J. M., editors, *The Neurobiology of Computation: Proceedings of the Annual Computation and Neural Systems Meeting*. Dordrecht, The Netherlands: Kluwer.
- Barrow, H. G., Bray, A. J., and Budd, J. M. L. (1996). A self-organizing model of ‘color blob’ formation. *Neural Computation*, 8:1427–1448.
- Bartfeld, E., and Grinvald, A. (1992). Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. *PNAS*, 89(24):11905–11909.
- Bartsch, A. P., and van Hemmen, J. L. (2001). Combined Hebbian development of geniculocortical and lateral connectivity in a model of primary visual cortex. *Biological Cybernetics*, 84:41–55.
- Bednar, J. A. (2002). *Learning to See: Genetic and Environmental Influences on Visual Development*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX. Technical Report AI-TR-02-294.
- Bednar, J. A., Kelkar, A., and Miikkulainen, R. (2002). Modeling large cortical networks with growing self-organizing maps. *Neurocomputing*, 44–46:315–321.

- Bednar, J. A., and Miikkulainen, R. (2000). Tilt aftereffects in a self-organizing model of the primary visual cortex. *Neural Computation*, 12(7):1721–1740.
- Bednar, J. A., and Miikkulainen, R. (2003). Self-organization of spatiotemporal receptive fields and laterally connected direction and orientation maps. *Neurocomputing*, 52–54:473–480.
- Bednar, J. A., and Miikkulainen, R. (2006). Joint maps for orientation, eye, and direction preference in a self-organizing model of V1. *Neurocomputing*, 69(10–12):1272–1276.
- Blakemore, C., and van Sluyters, R. C. (1975). Innate and environmental factors in the development of the kitten’s visual cortex. *The Journal of Physiology*, 248:663–716.
- Blasdel, G. G. (1992a). Differential imaging of ocular dominance columns and orientation selectivity in monkey striate cortex. *The Journal of Neuroscience*, 12:3115–3138.
- Blasdel, G. G. (1992b). Orientation selectivity, preference, and continuity in monkey striate cortex. *The Journal of Neuroscience*, 12:3139–3161.
- Blasdel, G. G., and Campbell, D. (2001). Functional retinotopy of monkey visual cortex. *The Journal of Neuroscience*, 21(20):8286–8301.
- Blasdel, G. G., and Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585.
- Bosking, W. H., Zhang, Y., Schofield, B. R., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *The Journal of Neuroscience*, 17(6):2112–2127.
- Bourgeois, J. P., Jastreboff, P. J., and Rakic, P. (1989). Synaptogenesis in visual cortex of normal and preterm monkeys: Evidence for intrinsic regulation of synaptic overproduction. *Proceedings of the National Academy of Sciences, USA*, 86:4297–4301.
- Brainard, D. H., Roorda, A., Yamauchi, Y., Calderone, J. B., Metha, A., Neitz, M., Neitz, J., Williams, D. R., and Jacobs, G. H. (2000). Functional consequences of the relative numbers of L and M cones. *Journal of the Optical Society of America A*, 17:607–614.
- Brown, P. K., and Wald, G. (1963). Visual pigments in human and monkey retinas. *Nature*, 200:37–43.

- Burger, T., and Lang, E. W. (1999). An incremental Hebbian learning model of the primary visual cortex with lateral plasticity and real input patterns. *Zeitschrift für Naturforschung C — A Journal of Biosciences*, 54:128–140.
- Burger, T., and Lang, E. W. (2001). Self-organization of local cortical circuits and cortical orientation maps: A nonlinear Hebbian model of the visual cortex with adaptive lateral couplings. *Zeitschrift für Naturforschung C—A Journal of Biosciences*, 56:464–478.
- Calkins, D. J. (2001). Seeing with S cones. *Progress in Retinal and Eye Research*, 20(3):255–287.
- Chance, F. S., Nelson, S. B., and Abbott, L. F. (1999). Complex cells as cortically amplified simple cells. *Nature Neuroscience*, 2(3):277–282.
- Chatterjee, S., and Callaway, E. M. (2002). S cone contributions to the magnocellular visual pathway in macaque monkey. *Neuron*, 35:1135–1146.
- Chatterjee, S., and Callaway, E. M. (2003). Parallel colour-opponent pathways to primary visual cortex. *Nature*, 426:668–670.
- Choe, Y. (2001). *Perceptual Grouping in a Self-Organizing Map of Spiking Neurons*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX. Technical Report AI01-292.
- Ciroux, J. (2005). *Simulating the McCollough Effect in a Self-Organizing Model of the Primary Visual Cortex*. Master’s thesis, The University of Edinburgh, Scotland, UK.
- Conway, B. R. (2001). Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). *Journal of Neuroscience*, 21(8):2768–2783.
- Coppola, D. M., White, L. E., Fitzpatrick, D., and Purves, D. (1998). Unequal representation of cardinal and oblique contours in ferret visual cortex. *Proceedings of the National Academy of Sciences, USA*, 95(5):2621–2623.
- Courtney, S. M., Finkel, L. H., and Buchsbaum, G. (1995). Network simulations of retinal and cortical contributions to color constancy. *Vision Research*, 35:413–434.
- Dacey, D. M. (1994). Physiology, morphology, and spatial densities of identified ganglion cell types in primate retina. In Bock, G. R., and Goode, J. A., editors, *Higher-Order Processing in the Visual System*, Ciba Foundation Symposium 184, 12–28. Hoboken, NJ: Wiley.

- Dacey, D. M. (2000). Parallel pathways for spectral coding in primate retina. *Annual Review of Neuroscience*, 23:743–775.
- De Paula, J. B. (2007a). Austin arboretum foliage corpora. Free for research use on condition you acknowledge the source. <http://nn.cs.utexas.edu/keyword?arboretum>.
- De Paula, J. B. (2007b). Flowers image corpus. Free for research use as long as you acknowledge the source. <http://nn.cs.utexas.edu/keyword?flowers>.
- DeAngelis, G. C., Ohzawa, I., and Freeman, R. D. (1995). Receptive-field dynamics in the central visual pathways. *Trends in Neurosciences*, 18:451–458.
- Delahunt, P. B., and Brainard, D. H. (2004). Does human color constancy incorporate the statistical regularity of natural daylight? *Journal of Vision*, 4:57–81.
- Durbin, R., and Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. *Nature*, 343:644–647.
- Einhauser, W., Kayser, C., Konig, P., and Kording, K. P. (2002). Learning the invariance properties of complex cells from their responses to natural stimuli. *European Journal of Neuroscience*, 15(3):475–486.
- Engel, S. A., and Furmanski, C. S. (2001). Selective adaptation to color contrast in human primary visual cortex. *Journal of Neuroscience*, 21(11):3949–3954.
- Fairchild, M. D., and Wyble, D. R. (1998). Colorimetric characterization of the Apple Studio display (flat panel LCD). Technical report, Munsell Color Science Laboratory, Rochester, NY.
- Gegenfurtner, K. R., and Kiper, D. C. (2003). Color vision. *Annual Review of Neuroscience*, 26:181–206.
- Geisler, W. S., Perry, J. S., Super, B. J., and Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41:711–724.
- Geisler, W. S., Thornton, T., Gallogly, D. P., and Perry, J. S. (2000). Image structure models of texture and contour visibility. In *Search and Target Acquisition (Recherche et acquisition d’objectifs)*, RTO Meeting Proceedings 45, 15/1–15/8. Hull, Québec: Canada Communication Group.
- Goodhill, G. J. (1993). Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics*, 69:109–118.

- Grinvald, A., Lieke, E. E., Frostig, R. D., and Hildesheim, R. (1994). Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. *The Journal of Neuroscience*, 14:2545–2568.
- Guiasu, S. (1977). *Information theory with applications*. New York: McGraw-Hill.
- Hata, Y., Tsumoto, T., Sato, H., Hagihara, K., and Tamura, H. (1993). Development of local horizontal interactions in cat visual cortex studied by cross-correlation analysis. *Journal of Neurophysiology*, 69:40–56.
- Haykin, S. (1994). *Neural Networks: A Comprehensive Foundation*. New York: Macmillan.
- Haykin, S. (1999). *Neural Networks, A Comprehensive Foundation*. Upper Saddle River, New Jersey: Prentice Hall.
- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. Hoboken, NJ: Wiley.
- Hirsch, H. V. B., and Spinelli, D. (1970). Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science*, 168:869–871.
- Hirsch, J. A., and Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat’s visual cortex. *The Journal of Neuroscience*, 11:1800–1809.
- Horn, D., Levy, N., and Ruppin, E. (1998). Memory maintenance via neuronal regulation. *Neural Computation*, 10:1–18.
- Hubel, D. H., and Livingstone, M. S. (1983). The 11th J.A.F. Stevenson memorial lecture: Blobs and color vision. *Canadian Journal of Physiology and Pharmacology*, 61(12):1433–1441.
- Hubel, D. H., and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of Physiology*, 160:106–154.
- Hubel, D. H., and Wiesel, T. N. (1963). Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *Journal of Neurophysiology*, 26(6):994–1002.
- Hubel, D. H., and Wiesel, T. N. (1967). Cortical and callosal connections concerned with the vertical meridian of visual fields of the cat. *Journal of Neurophysiology*, 30:1561–1573.

- Hubel, D. H., Wiesel, T. N., and LeVay, S. (1977). Plasticity of ocular dominance columns in monkey striate cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 278:377–409.
- Hubener, M., and Bolz, J. (1992). Relationships between dendritic morphology and cytochrome oxidase compartments in monkey striate cortex. *Journal of Comparative Neurology*, 324:67–80.
- Issa, N. P., Trepel, C., and Stryker, M. P. (2001). Spatial frequency maps in cat visual cortex. *The Journal of Neuroscience*, 20:8504–8514.
- Itten, J. (1970). *The elements of color*. New York, NY: Wiley.
- Jacobs, G. H., Williams, G. A., Cahill, H., and Nathans, J. (2007). Emergence of novel color vision in mice engineered to express a human cone photopigment. *Science*, 315:1723–1725.
- Jordan, G., and Mollon, J. D. (1993). A study of women heterozygous for colour deficiencies. *Vision Research*, 33:1495–1508.
- Kalarickal, G. J., and Marshall, J. A. (1998). Modeling dynamic receptive field changes in area V1 using inhibitory learning. In Bower, J. M., editor, *Computational Neuroscience: Trends in Research, 1998*. New York: Plenum Press.
- Kandel, E. R., Schwartz, J. H., and Jessell, T. M. (2000). *Principles of Neural Science*. New York: McGraw-Hill. Fourth edition.
- Kee, F. (2005). God spilled the paint. Copyright Frank Kee. Desert Flowers on the east side of the Carrizo plain, in the Temblor Range, about 50 miles due west of Bakersfield. <http://www.keesphotos.com> (accessed Sun Jun 24 17:17:49 CDT 2007).
- Kingdom, F. A. (2003). Color brings relief to human vision. *Nature Neuroscience*, 6(6):641–644.
- Kingdom, F. A., Rangwala, S., and Hammamji, K. (2005). Chromatic properties of the colour-shading effect. *Vision Research*, 45:1425–1437.
- Klug, K., Herr, S., Ngo, I. T., Sterling, P., and Schein, S. (2003). Macaque retina contains an S-cone OFF midget pathway. *Journal of Neuroscience*, 23:9881–9887.
- Kohonen, T. (1982a). Analysis of a simple self-organizing process. *Biological Cybernetics*, 44:135–140.

- Kohonen, T. (1982b). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69.
- Kohonen, T. (2001). *Self-Organizing Maps*. Berlin: Springer. Third edition.
- Landisman, C. E., and Ts'o, D. Y. (2002a). Color processing in macaque striate cortex: Electrophysiological properties. *Journal of Neurophysiology*, 87(6):3138–3151.
- Landisman, C. E., and Ts'o, D. Y. (2002b). Color processing in macaque striate cortex: Relationships to ocular dominance, cytochrome oxidase, and orientation. *Journal of Neurophysiology*, 87(6):3126–3137.
- Law, J. S., and Bednar, J. A. (2006). Surround modulation by long-range lateral connections in an orientation map model of primary visual cortex development and function. Abstract presented at the 2006 Meeting of Society for Neuroscience. Atlanta, GA. (546.4/U14).
- Lee, B. B., Kremers, J., and Yeh, T. (1998). Receptive fields of primate retinal ganglion cells studied with a novel technique. *Visual Neuroscience*, 15:161–175.
- Lee, C. W., Eglén, S. J., and Wong, R. O. L. (2002a). Segregation of ON and OFF retinogeniculate connectivity directed by patterned spontaneous activity. *Journal of Neurophysiology*, 88(5):2311–2321.
- Lee, T.-W., Wachtler, T., and Sejnowski, T. J. (2002b). Color opponency is an efficient representation of spectral properties in natural scenes. *Vision Research*, 42:2095–2103.
- Lennie, P. (2000). Color vision: Putting it together. *Current Biology*, 10:R589–591.
- Levitt, J. B., Kiper, D. C., and Movshon, J. A. (1994). Receptive fields and functional architecture of macaque V2. *JNP*, 71:2517–2542.
- Livingstone, M. S., and Hubel, D. H. (1984). Anatomy and physiology of a color system in the primate visual cortex. *The Journal of Neuroscience*, 4(1):309–356.
- Lovell, P. G., Tolhurst, D. J., Párraga, C. A., Baddeley, R. J., Leonards, U., and Troscianko, J. (2005). Stability of the color-opponent signals under changes of illuminant in natural scenes. *Journal of the Optical Society of America A—Optics Image Science and vision*, 22:2060–2071.

- Löwel, S. (1994). Ocular dominance column development: Strabismus changes the spacing of adjacent columns in cat visual cortex. *The Journal of Neuroscience*, 14(12):7451–7468.
- Löwel, S., and Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255:209–212.
- Mansfield, R. J. W. (1974). Neural basis of orientation perception in primate vision. *Science*, 186:1133–1135.
- McCollough, C. (1965). Color adaptation of edge-detectors in the human visual system. *Science*, 149(3688):1115–1116.
- Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (2005). *Computational Maps in the Visual Cortex*. Berlin: Springer.
- Miller, K. D. (1994). 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. *The Journal of Neuroscience*, 14:409–441.
- Miller, K. D., Erwin, E., and Kayser, A. (1999). Is the development of orientation selectivity instructed by activity?. *Journal of Neurobiology*, 41:44–57.
- Miranker, W. L. (1997). Consciousness is an information field induced by Hebbian dynamics. TICAM Report 97-11, Texas Institute for Computational and Applied Mathematics, The University of Texas at Austin, Austin, TX.
- Nakagama, H., and Tanaka, S. (2004). Self-organization model of cytochrome oxidase blobs and ocular dominance columns in the primary visual cortex. *Cerebral Cortex*, 14(4):376–386.
- Nascimento, S. M. C., and Ferreira, F. P. (2002). Statistics of spatial cone-excitation ratios in natural scenes. *Journal of the Optical Society of America A.*, 19:1484–1490.
- Obermayer, K., Ritter, H., and Schulten, K. J. (1990a). Large-scale simulations of self-organizing neural networks on parallel computers: Application to biological modelling. *Parallel Computing*, 14:381–404.
- Obermayer, K., Ritter, H., and Schulten, K. J. (1990b). A principle for the formation of the spatial structure of cortical feature maps. *Proceedings of the National Academy of Sciences, USA*, 87:8345–8349.

- Obermayer, K., Sejnowski, T. J., and Blasdel, G. G. (1995). Neural pattern formation via a competitive Hebbian mechanism. *Behavioural Brain Research*, 66(1-2):161–167.
- Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15:267–273.
- Oja, E. (1991a). Neural networks in image processing and analysis. Technical report, Department of Computer Science, Tokyo Institute of Technology, Tokyo, Japan.
- Oja, E. (1991b). PCA neural networks. In *Proceedings of the International Symposium MTNS 91, Kobe, Japan*.
- Olmos, A., and Kingdom, F. A. A. (2004). McGill calibrated colour image database. <http://tabby.vision.mcgill.ca>.
- Olshausen, B. A., and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609.
- Packer, O. S., and Dacey, D. M. (2002). Receptive field structure of H1 horizontal cells in macaque monkey retina. *Journal of Vision*, 2:272–292.
- Pallas, S. L., and Finlay, B. L. (1991). Compensation for population-size mismatches in the hamster retinotectal system: Alterations in the organization of retinal projections. *Visual Neuroscience*, 6:271–281.
- Palmer, C. M., and Bednar, J. A. (2006). Modeling the development of topographic and laminar organization for orientation and spatial frequency in the primary visual cortex. In *Society for Neuroscience Abstracts*. Society for Neuroscience, www.sfn.org. Program No. 546.3.
- Parraga, C. A., Brelstaff, G., Troscianko, T., and Moorhead, I. R. (1998). Color and luminance information in natural scenes. *Journal of the Optical Society of America A*, 15:563–569.
- Piepenbrock, C., Ritter, H., and Obermayer, K. (1997). The joint development of orientation and ocular dominance: Role of constraints. *Neural Computation*, 9:959–970.
- Ramtohl, T. (2006). *A Self-Organizing Model of Disparity Maps in the Primary Visual Cortex*. Master’s thesis, The University of Edinburgh, Scotland, UK.

- Rao, A. R., Cecchi, G., Peck, C., and Kozloski, J. (2005). Self-organizing cortical color maps. In *Proceedings of the 17th Annual Symposium of the International Society of Optical Engineering on Electronic Imaging*. San Jose, CA.
- Rodieck, R. W. (1965). Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vision Research*, 5(11):583–601.
- Roe, A. W., and Ts'o, D. Y. (1995). Visual topography in primate V2: Multiple representation across functional stripes. *Journal of Neuroscience*, 15:3689–3715.
- Roe, A. W., and Ts'o, D. Y. (1999). Specificity of color connectivity between primate V1 and V2. *Journal of Neurophysiology*, 82:2719–2730.
- Rosa, M. G. P., Fritsches, K. A., and Elston, G. N. (1997). The second visual area in the marmoset monkey: Visuotopic organisation, magnification factors, architectonical boundaries, and modularity. *Journal of Comparative Neurology*, 387:547–567.
- Rowe, M. H. (2002). Trichromatic color vision in primates. *News in Physiological Sciences*, 17(3):93–98.
- Saarinen, J., and Kohonen, T. (1985). Self-organized formation of colour maps in a model cortex. *Perception*, 14:711–719.
- Sato, H., Katsuyama, N., Tamura, H., Hata, Y., and Tsumoto, T. (1994). Broad-tuned chromatic inputs to color-selective neurons in the monkey visual cortex. *Journal Neurophysiology*, 72:163–168.
- Seidemann, E., Poirson, A. B., Wandell, B. A., and Newsome, W. T. (1999). Color signals in area MT of the macaque monkey. *Neuron*, 24:911–917.
- Sharpe, L. T., Stockman, A., Jagle, H., and Nathans, J. (1999). Opsin genes, cone photopigments, color vision, and color blindness. In Gegenfurtner, K., and Sharpe, L. T., editors, *Color vision: From genes to perception*, 3–51. Cambridge, UK: Cambridge University Press.
- Shipp, S., and Zeki, S. (2002a). The functional organization of area V2, I: Specialization across stripes and layers. *Visual Neuroscience*, 19:187–210.
- Shipp, S., and Zeki, S. (2002b). The functional organization of area V2, II: The impact of stripes on visual topography. *Visual Neuroscience*, 19:211–231.
- Simoncelli, E. P., and Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24:1193–1216.

- Sincich, L. C., and Blasdel, G. G. (2001). Oriented axon projections in primary visual cortex of the monkey. *The Journal of Neuroscience*, 21:4416–4426.
- Sincich, L. C., and Horton, J. C. (2002a). Divided by cytochrome oxidase: A map of the projections from V1 to V2 in macaques. *Science*, 295:1734–1737.
- Sincich, L. C., and Horton, J. C. (2002b). Pale cytochrome oxidase stripes in V2 receive the richest projections from macaque striate cortex. *Journal of Comparative Neurology*, 447:18–33.
- Sincich, L. C., and Horton, J. C. (2005a). The circuitry of V1 and V2: Integration of color, form, and motion. *ARN*, 28:303–326.
- Sincich, L. C., and Horton, J. C. (2005b). Input to V2 thin stripes arises from V1 cytochrome oxidase patches. *JN*, 24:10087–10093.
- Sirosh, J. (1995). *A Self-Organizing Neural Network Model of the Primary Visual Cortex*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX. Technical Report AI95-237.
- Sirosh, J., and Miikkulainen, R. (1994). Cooperative self-organization of afferent and lateral connections in cortical maps. *Biological Cybernetics*, 71:66–78.
- Sirosh, J., Miikkulainen, R., and Bednar, J. A. (1996). Self-organization of orientation maps, lateral connections, and dynamic receptive fields in the primary visual cortex. In Sirosh, J., Miikkulainen, R., and Choe, Y., editors, *Lateral Interactions in the Cortex: Structure and Function*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://nn.cs.utexas.edu/web-pubs/htmlbook96>.
- Sit, Y. F., and Miikkulainen, R. (2006). Self-organization of hierarchical visual maps with feedback connections. *Neurocomputing*, 69:1309–1312.
- Sit, Y. F., and Miikkulainen, R. (2007). Computational predictions on the receptive fields and organization of V2 for shape processing. Submitted to *Neural Computation*.
- Stanikunas, R., Vaitkevicius, H., and Kulikowski, J. J. (2004). Investigation of color constancy with a neural network. *Neural Networks*, 17:327–337.
- Stockman, A., and Sharpe, L. T. (2000). The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research*, 40:1711–1737.

- Stockman, A., Sharpe, L. T., and Fach, C. (1999). The spectral sensitivity of the human short-wavelength sensitive cones derived from thresholds and color matches. *Vision Research*, 39:2901–2927.
- Sugita, Y. (2004). Experience in early infancy is indispensable for color perception. *Current Biology*, 14(14):1267–1271.
- Switkes, E., Mayer, M. J., and Sloan, J. A. (1978). Spatial frequency analysis of the visual environment: Anisotropy and the carpentered environment hypothesis. *Vision Research*, 18(10):1393–1399.
- Taylor, D. R., Finkel, L. H., and Buchsbaum, G. (2000). Color-opponent receptive fields derived from independent component analysis of natural images. *Vision Research*, 40:2671–2676.
- Tootell, R. B. H., Silverman, M. S., and De Valois, R. L. (1981). Spatial frequency columns in primary visual cortex. *Science*, 214:813–815.
- Tootell, R. B. H., Silverman, M. S., Hamilton, S. L., Switkes, E., and De Valois, R. L. (1988). Functional anatomy of macaque striate cortex. V. Spatial frequency. *The Journal of Neuroscience*, 8:1610–1624.
- Troscianko, T., Párraga, C. A., Leonards, U., Baddeley, R. J., Troscianko, J., and Tolhurst, D. J. (2003). Leaves, fruit, shadows, and lighting in Kibale forest, Uganda. *Perception*, 32 Supplement.
- Tsien, J. Z. (2000). Linking Hebb’s coincidence-detection to memory formation. *Current Opinion in Neurobiology*, 10(2):266–273.
- Ts’o, D. Y., Frostig, R. D., Lieke, E. E., and Grinvald, A. (1990). Functional organization of primate visual cortex revealed by high resolution optical imaging. *Science*, 249:417–420.
- Turrigiano, G. G. (1999). Homeostatic plasticity in neuronal networks: The more things change, the more they stay the same. *Trends in Neurosciences*, 22(5):221–227.
- Usui, S., Nakauchi, S., and Miyake, S. (1994). Acquisition of color opponent representation by a three-layered neural network model. *Biological Cybernetics*, 72:35–41.
- Usui, S., Nakauchi, S., and Miyamoto, Y. (1992). A decorrelating neural network model for color constancy. *Proceedings of the international conference of the IEEE*, 14:1030–1031.

- van Hateren, J. H., and Ruderman, D. L. (1998). Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proceedings of the Royal Society London B.*, 265:2315–2320.
- van Hateren, J. H., and van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society London B.*, 265:359–366.
- Vanduffel, W., Tootell, R. B., Schoups, A. A., and Orban, G. A. (2002). The organization of orientation selectivity throughout macaque visual cortex. *Cerebral Cortex*, 12:647–662.
- Visual Delights (1999). Austin and vicinity. CD of images for sale by company. Copyright Visual Delights, Inc., all rights reserved. <http://www.visualdelights.net>.
- von der Malsburg, C. (1973). Self-organization of orientation-sensitive cells in the striate cortex. *Kybernetik*, 15:85–100. Reprinted in Anderson and Rosenfeld (1988), 212–227.
- Vos, J. J. (1978). Colorimetric and photometric properties of a 2-degree fundametal observer. *Color research and application*, 3:125–128.
- Webster, M. A., and Mollon, J. D. (1999). Adaptation and the color statistics of natural images. *Vision Research*, 37(23):3283–3298.
- Weliky, M., Kandler, K., Fitzpatrick, D., and Katz, L. C. (1995). Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron*, 15:541–552.
- Wilson, T. (2004). Personal communication. Spectral Power Distribution function for a Sony GDM 500 CRT monitor, provided April 27.
- Wong, R. O. L. (1999). Retinal waves and visual system development. *Annual Review of Neuroscience*, 22:29–47.
- Xiao, Y., Casti, A., Xiao, J., and Kaplan, E. (2007). Hue maps in primate striate cortex. *NeuroImage*, 35:771–786.
- Xiao, Y., and Fellerman, D. J. (2004). Projections from primary visual cortex to cytochrome oxidase thin stripes and interstripes of macaque visual area 2. *PNAS*, 101:7147–7151.

- Xiao, Y., Wang, Y., and Felleman, D. J. (2003). A spatially organized representation of color in macaque cortical area V2. *Nature*, 421:535–539.
- Yao, Y. Y. (1993). Information-theoretic measures for knowledge discovery and data mining. In Karmeshu, editor, *Entropy measures, maximum entropy principle and emerging applications. Studies in fuzziness and soft computing, Volume 119*, 115–136. Berlin: Springer.
- Yoshioka, T., and Dow, B. M. (1996). Color, orientation and cytochrome oxidase reactivity in areas V1, V2, and V4 of macaque monkey visual cortex. *Behavioural Brain Research*, 76:71–88.

Vita

Judah De Paula was born in Fontana, California on March 19, 1978, the first son of Bruce and Audrey De Paula. Over a period of five years he attended San Bernardino Valley Community College, Kennesaw State University, and the Georgia Institute of Technology; at the end receiving a Bachelor of Science in Computer Science with a minor and a certificate in Cognitive Sciences. In the Fall of 2000 he entered graduate school at the University of Texas at Austin, and in 2006 he married Brooke, a native Austinite. He currently works at Enthought, a scientific computing software company.

Permanent Address: 12304 Abney Dr. Apt A.
Austin, Texas 78729

This dissertation was typeset with $\text{\LaTeX} 2_{\epsilon}$ ¹ by the author.

¹ $\text{\LaTeX} 2_{\epsilon}$ is an extension of \LaTeX . \LaTeX is a collection of macros for \TeX . \TeX is a trademark of the American Mathematical Society. The macros used in formatting this dissertation were written by Dinesh Das, Department of Computer Sciences, The University of Texas at Austin, and extended by Bert Kay, James A. Bednar, Ayman El-Khashab, and Judah De Paula.