THE UNIVERSITY of EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

Modeling the emergence of perceptual color space in the primary visual cortex

Christopher E. Ball

Doctor of Philosophy
Institute for Adaptive and Neural Computation
School of Informatics
University of Edinburgh

2014

Abstract

Humans' perceptual experience of color is very different from what one might expect, given the light reaching the eye. Identical patterns of light are often perceived as different colors, and different patterns of light are often perceived as the same color. Even more strikingly, our perceptual experience is that hues are arranged circularly (with red similar to violet), even though single-wavelength lights giving rise to perceptions of red and violet are at opposite ends of the wavelength spectrum. The goal of this thesis is to understand how perceptual color space arises in the brain, focusing on the arrangement of hue. To do this, we use computational modeling to integrate findings about light, physiology of the visual system, and color representation in the brain.

Recent experimental work shows that alongside spatially contiguous orientation preference maps, macaque primary visual cortex (V1) represents color in isolated patches, and within those patches hue appears to be spatially organized according to perceptual color space. We construct a model of the early visual system that develops based on natural input, and we demonstrate that several factors interact to prevent this first model from developing a realistic representation of hue. We show these factors as independent dimensions and relate them to problems the brain must be overcoming in building a representation of perceptual color space: physiological and environmental variabilities to which the brain is relatively insensitive (surprisingly, given the importance of input in driving development). We subsequently show that a model with a certain position on each dimension develops a hue representation matching the range and spatial organization found in macaque V1—the first time a model has done so. We also show that the realistic results are part of a spectrum of possible results, indicating other organizations of color and orientation that could be found in animals, depending on physiological and environmental factors. Finally, by analyzing how the models work, we hypothesize that well-accepted biological mechanisms such as adaptation, typically omitted from models of both luminance and color processing, can allow the models to overcome these variabilities, as the brain does.

These results help understand how V1 can develop a stable, consistent representation of color despite variabilities in the underlying physiology and input statistics. This in turn suggests how the brain can build useful, stable representations in general based on visual experience, despite irrelevant variabilities in input and physiology. The resulting models form a platform to investigate various adult color visual phenomena, as well as to predict results of rearing experiments.

Acknowledgements

To the memory of Rick Jones
1953–2014

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Christopher E. Ball)

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Color vision

When we look at a rainbow, we see a spectrum of colors. The sun's light has been dispersed by water droplets, and the colors seen vary by wavelength. In a simpler configuration, shining white light—i.e. light with a flat spectral power distribution—into a prism results in a similar spectrum of colors (figure 1.1(a)). A number of features of this simple and familiar spectrum are, however, puzzling:

- We perceive bands of color, with varying widths and varying brightnesses—but none of these three phenomena matches the spectral distribution of solar power (figure 1.1(c)).

- We can see that red and violet are at opposite ends of the spectrum, yet in a space that describes human judgments of perceptual similarity between colors, red and violet are adjacent (figure 1.1(b))

- Some colors do not appear in the spectrum at all. For example, colors between red and violet (such as magenta) can only appear using a mixture of wavelengths, yet these seem just as colorful and distinct as those experienced for monochromatic light.

These observations indicate that the colors we perceive (perceptual color space) and the light reaching our eyes at any point (wavelength space) are not as clearly related as we might imagine.

Beyond the spectrum of colors we see from the dispersion of white light, there are many other phenomena that illustrate how the light entering our eyes at a particular

1
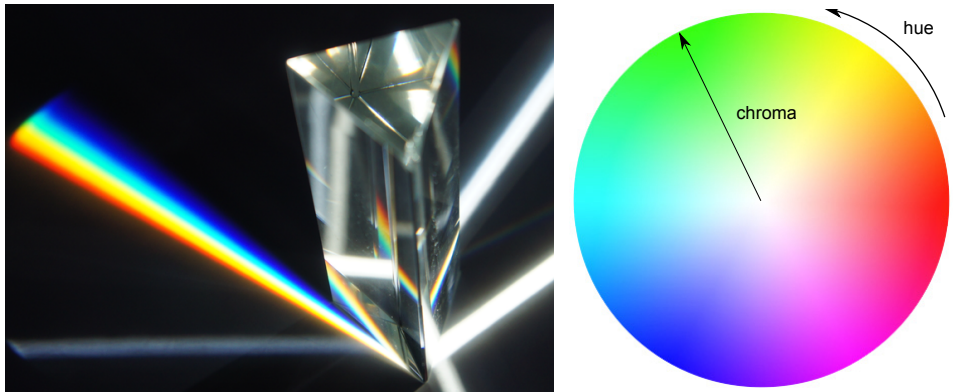
(a) Sunlight dispersed by a prism

(b) Perceptual hue space
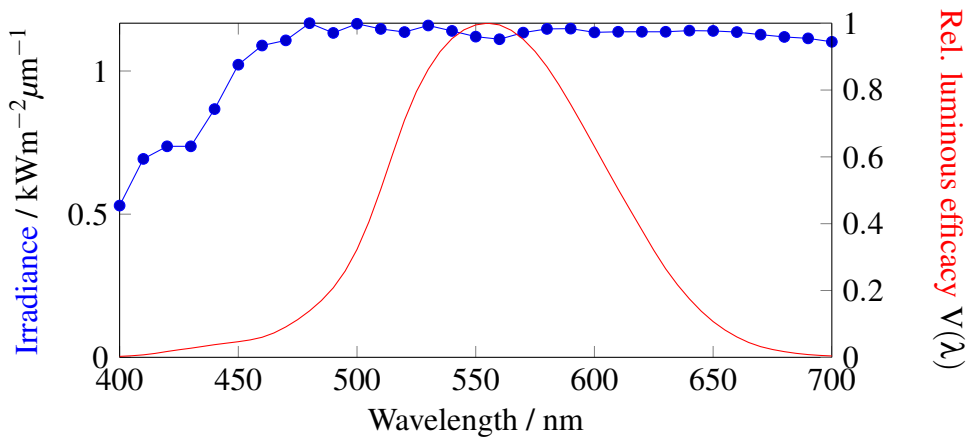
(c) Sunlight's spectral power distribution, and human relative sensitivity

Figure 1.1: *Perception of color. Wavelength space is the spectral power distribution of light coming to the eye. In (a), white light from the sun is dispersed according to wavelength by a prism. We perceive bands of color (of varying widths and perceived brightnesses), and we judge long- and short-wavelength light as being similar to each other, as shown by the joining of red and blue in (b). (b) perceptual color space: hue arranged in a circle so that similar hues are close together, with chroma shown radially. A third dimension of color vision, lightness, is not shown. Approximately circular spaces like this can be constructed from human judgments of similarities of pairs of hues. (c) The blue line plots solar irradiance at the Earth's surface, showing power is reasonably constant over much of the visible range of wavelengths. The red line plots human relative sensitivity as a function of wavelength, showing we are typically most sensitive to mid-spectrum light (in daylight). A 550 nm light typically appears yellow, and for a given physical power will appear much brighter than e.g. 650 nm ("red") light. Note: colors shown in (a) and (b) are only approximate. Solar irradiance and $V(\lambda)$ data from Wyszecki and Stiles (1982).*

location does not predict the color we perceive. For instance, lights of different spectral power distributions viewed in isolation can give rise to the same perceived color (metamerism; Brainard 2003), or increasing only the luminance of a stimulus can change its perceived color (e.g. even a single wavelength of light appearing red can instead appear yellow as luminance is increased—the Bezold-Brücke effect; Purdy, 1931).

The phenomena mentioned so far have been for relatively simple configurations of light, unlike a typical natural scene, which contains surfaces as well as an illuminant. When we consider natural scenes, a number of other interesting observations arise. For example:

- Different spectral power distributions reaching the eye at a point can appear as the same color, depending on surroundings. Figure 1.2 shows that we have a degree of color constancy, here demonstrated by our ability to "discount the illuminant". This ability to separate the underlying reflectance of an object—a fixed property of the object—from the illuminant is something we take for granted, but is a key feature of color vision.

- On the other hand, a physically identical object presented in different contexts can appear different colors, even if the light it reflects to our eyes is the same in all cases (figures 1.3(a) and 1.3(b)). The context could be spatial—i.e. the object's color may appear to change because of surrounding colors—or the context could be temporal—i.e. the object's color may appear to change because of colors previously viewed (figure 1.4).

Clearly, the color vision system is not attempting to reconstruct or distinguish spectral power distributions from the light reaching the eye. So, what is its goal? We consider this question in the next section.

1.2   Purpose of color vision

Under low-power light, a lump of coal appears black and a sheet of paper appears white. Under high-power light, the coal still appears black, and the paper still appears white. However, in the second situation the amount of light coming into the eye (the radiance) from the coal is higher than the radiance from the paper in the first situation. Our visual system is more concerned with how an object reflects light *relative* to other objects than it is with radiance. Reflected light from an object, relative to its surroundings, gives information about the object itself, rather than about the illuminant. We are
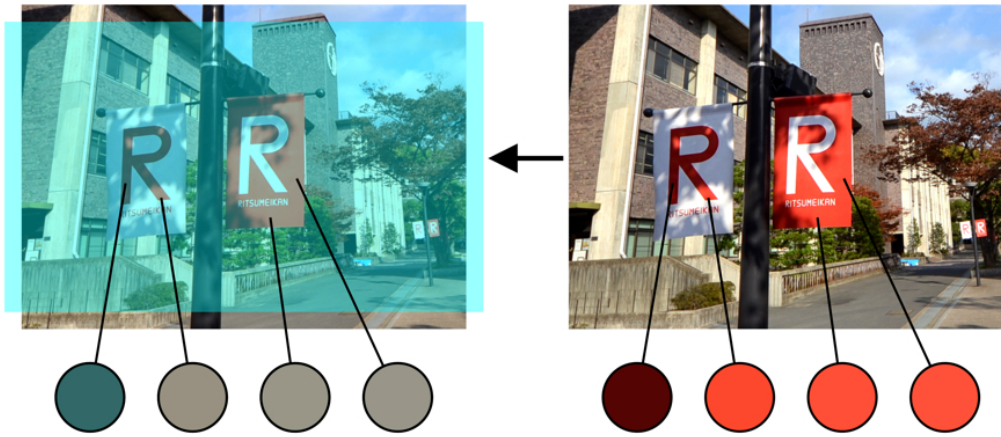
Figure 1.2: *Chromatic adaptation: "discounting the illuminant". The right-hand flag appears reddish in both images, despite being entirely gray in the left image (the circles beneath show the physical reality). This is simulating our adaptation to the illuminant, and is an important feature of our vision, helping to recognize objects despite changes in illuminant. Images reprinted from Kitaoka (2012).*

looking for things to eat, or things that will eat us; we need to be able to distinguish objects regardless of viewing conditions. The example above is for lightness, but the situation is similar for color. Whether an object is lit by midday sun or golden-hour light before sunset, reflecting much more long-wavelength light to our eyes in the second case, we still want to recognize the object (figure 1.5). Similarly, in a particular scene, the light returned from one surface may vary greatly from one end to the other, but the surface will usually still appear as one color, helping us to recognize the surface.

In addition to irrelevant external variety, our visual system must also deal with irrelevant internal variety. For instance, our eye's color equipment varies across the retina (chapter 2), but we need to recognize objects as the same wherever they occur in our field of view. And, as we age, our lens becomes increasingly yellow, but again we need stable identification of objects over time, rather than to continually perceive that our lens is changing.

We have seen how perceptual color space—with its separation of color into dimensions including chroma and circular hue—differs from wavelength space. We have also seen that color vision involves comparisons across space and time to help recognize the underlying properties of objects, giving us a degree of "color constancy". These are two key features of color vision, and there has been much research into both, but we still do not understand clearly how either occurs. This thesis aims to contribute towards understanding the first aspect: how does the transformation from wavelength space to
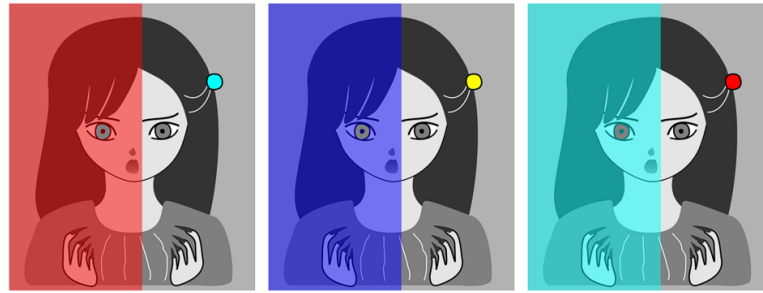
(a) All six eyes are physically identical (gray), but the eyes on the left appear colored due to the surrounding hues.
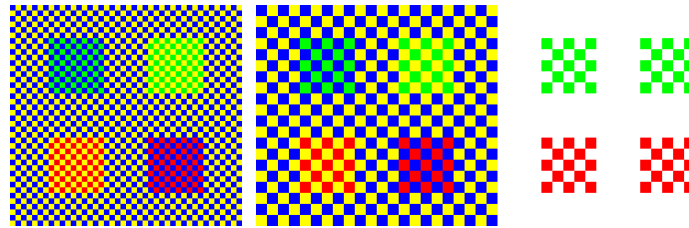
(b) The same four components ("green", "red", "yellow", and "blue") are used in all three images. On the right, green and red squares are shown in isolation; these squares are physically the same in the middle image but are surrounded by blue and yellow, and appear to have changed color. On the left, the smallest squares are still the same four colors, but are smaller, making the effect even stronger.

Figure 1.3: *Spatial chromatic adaptation.   (a) An example of chromatic induction (simultaneous contrast): the perceived color is shifted away from the inducer. Images reprinted from Kitaoka (2013). (b) Another example of chromatic induction (assimilation): the perceived color is shifted towards the inducer. The above phenomena demonstrate that color depends on more than the physical stimulus at one location.*
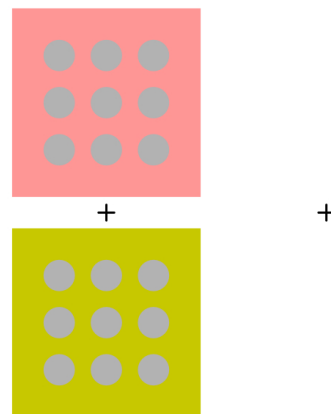
Figure 1.4: *Temporal chromatic adaptation. Fixate on the left-hand cross for 30 seconds, then look at the right-hand cross; you will see colored circles on backgrounds colored differently from the original image. The color we perceive depends on more than the physical stimulus at the current time. Images reprinted from Kitaoka (2008).*
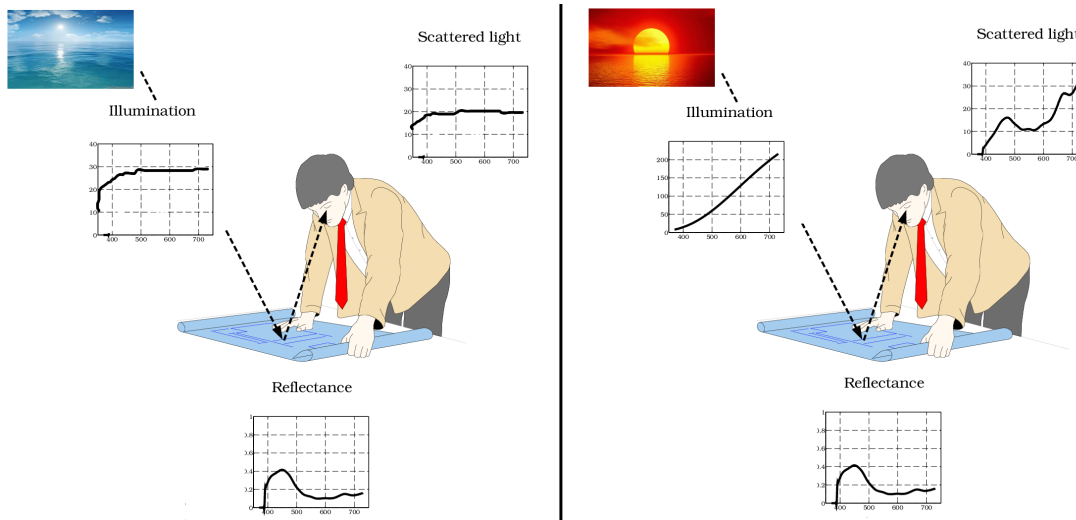
Figure 1.5: *Reflectance vs. radiance. Illumination from a source such as the sun, after it has interacted with the Earth's atmosphere and reaches the surface, has a spectral power distribution (see figure 1.1(c)) that can vary dramatically. We see an object when it reflects light to our eyes; the spectral reflectance of an object is a property of the object itself, and describes the relative fraction of light it reflects as a function of wavelength. However, the spectral power distribution of light actually returned to our eyes ("scattered" light in the illustration, a product of the illuminant and reflectance spectra) varies greatly because of changes in the illuminant. Our visual system is more concerned with information about objects than about illuminants: a blue object is perceived in both these cases, despite the very different light reaching the eyes. Adapted from Wandell (1995).*

perceptual color space happen in the early visual system? Understanding perceptual color space is an important step on the way to understanding color vision, and for understanding how our perceptual experiences in general are constructed by our brain.

Before we proceed to the next section, a note to the reader: we will be covering several different fields in the following chapters, so a glossary of terms is given in appendix B, along with a list of acronyms in appendix A.

1.3   Explaining perceptual color space

In this thesis we will investigate how the transformation from wavelength space (WS) to perceptual color space (PCS) could occur in the early visual system. We will argue below that complete genetic specification of PCS seems unlikely, while at the same time it is not obvious what physical characteristics of light reaching the eye would explain important features of PCS such as hue being circular (Lotto and Purves, 2002).

Light $\longrightarrow$ Eye $\longrightarrow$ Brain $\longrightarrow$ Behavior
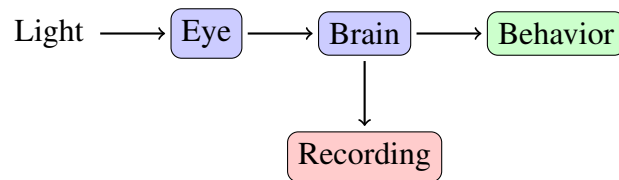
Recording

Figure 1.6: *Learning about the representation of perceptual experience in the brain. We can learn about the representation of color in the brain by e.g. presenting stimuli, and then recording activity in neurons in the brain (typically in laboratory animals, using various methods, such as with micro-electrodes, or non-invasively with various imaging techniques). We can also learn about the representation of color in the brain through behavioral outputs of humans (or other animals). For instance, we could test whether humans are able to distinguish various pairs of stimuli, such as two different wavelengths of light, or two of the physically identical eyes in figure 1.3(a).*

The transformation must happen in the brain, starting from light first reaching the photoreceptors. What is the physical mechanism that could account for the transformation? Furthermore, what factors combine to drive the development of such a mechanism? That is, why should we have this particular mechanism, as opposed to other possible color-processing machinery? This thesis aims to contribute towards answering these questions. We begin by considering where in the brain PCS is represented, since this will at least give us a bound on the mechanisms involved.

1.3.1 Where does perceptual color space arise?

By "where PCS is represented in the brain", we mean a representation of color stimuli that is organized in a way that reflects aspects of PCS. One of the main aims of neuroscience is to understand how neurons represent our perceptual experiences, and how those perceptions give useful information about the world (e.g. Zaidi et al. 2013). As illustrated in figure 1.6, we can do this by (a) asking humans to perform tasks, such as testing if they perceive two colors to be different (i.e. learn about perception based on behavior), and (b) recording patterns of brain activity in humans and animals when they are observing particular stimuli.

Through techniques such as those mentioned above, combined with anatomical studies, we have an understanding of the primate early visual system. The basic pathway of information flow from the eyes to visual cortex is shown in figure 1.7. We single out visual cortex as a critical stage in the processing of visual input because it has been found to represent features of the input in a particularly interesting way. We know, for instance, that there are neurons tuned to various orientations, such that primary visual
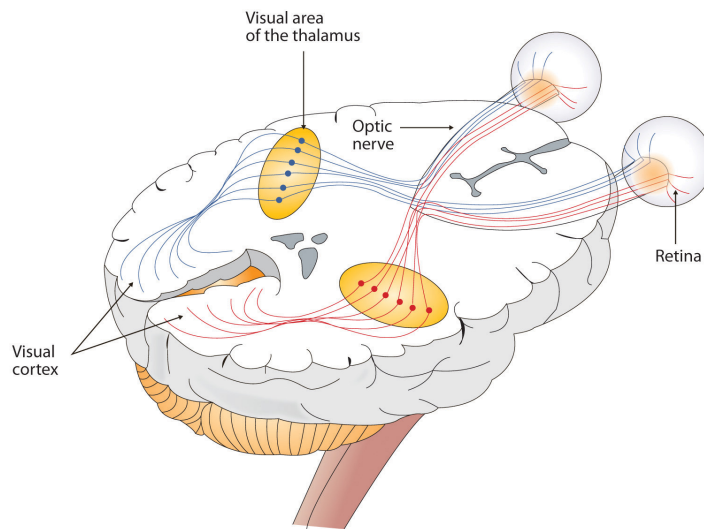
Figure 1.7: *Early visual system. Light is transduced to neural signals by photoreceptors at the back of the eye, on the retina. These signals then travel through the brain, via the lateral geniculate nucleus of the thalamus, to the visual cortex at the back and near the surface of the brain. The surface of the visual cortex has a spatially organized representation of many features of visual input (such as orientation and ocular dominance). There is evidence that color is also spatially organized according to perceptual color. Image reprinted from Strangor (2010).*

cortex (V1) contains a spatially contiguous map for orientation preference, organized retinotopically (Blasdel, 1992b; figure 1.8). A small area of map represents all orientations possible at a given location on the retina, and orientation preference generally varies smoothly. Additionally, the preferences are robust to various changes in the input, such as changes in contrast (Sclar and Freeman, 1982). Coding for orientation based on location in the map, i.e. a place code, is a reasonable assumption, because the pattern of neural activity will give subsequent neural architecture a reliable indication of the input orientation at each retinal location. While it is not the only coding possible, this kind of coding is simple, and is assumed by a wide body of neuroscience (Arbib, 2002). There are co-located maps of various features, including ocular dominance, motion direction, disparity, and spatial frequency, so it seems reasonable to look for a similar map for color.

Identifying whether there is such an organization for color—and if so, where it is—has been the subject of much research, but a definitive location has not yet been identified. Many studies involve macaque monkey, because it has a trichromatic visual system similar to our own, and shows similar performance on color vision tasks (chapter 2). In the following chapter, we will break up the macaque early visual system into the following key stages (illustrated in figure 1.9), and consider whether they show an
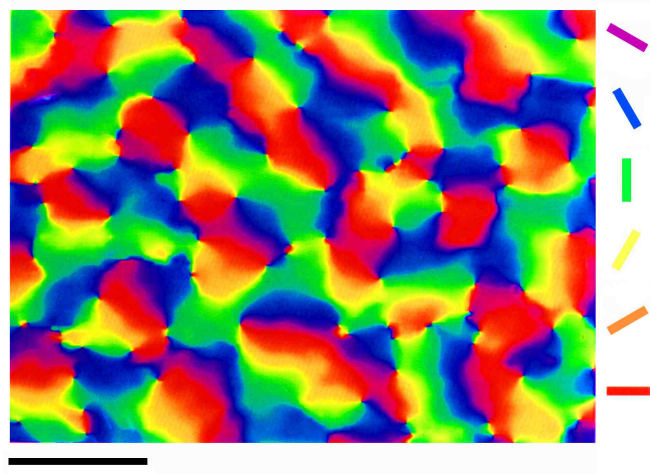
Figure 1.8: *Coding of orientation in visual cortex. Optical imaging of macaque monkey V1 reveals neurons with preferences for edges of different orientations. A monkey views a grating pattern, the orientation of which is varied while the activities of neurons in visual cortex are recorded by an optical camera. Each pixel is color-coded according to the orientation that caused the strongest response at that location. The map is smoothed both by the resolution of the camera, and in post-processing; an individual pixel likely represents around 1000 neurons. Similar orientation preferences tend to be nearby (the map is generally smooth), and many orientations are represented in a small area (giving good coverage). Orientation maps in macaque V1 are well accepted, but an equivalent representation for color has been more elusive. Scale bar:* 1 mm. *Image reprinted from Swindale (2008); data from Blasdel and Salama (1986).*
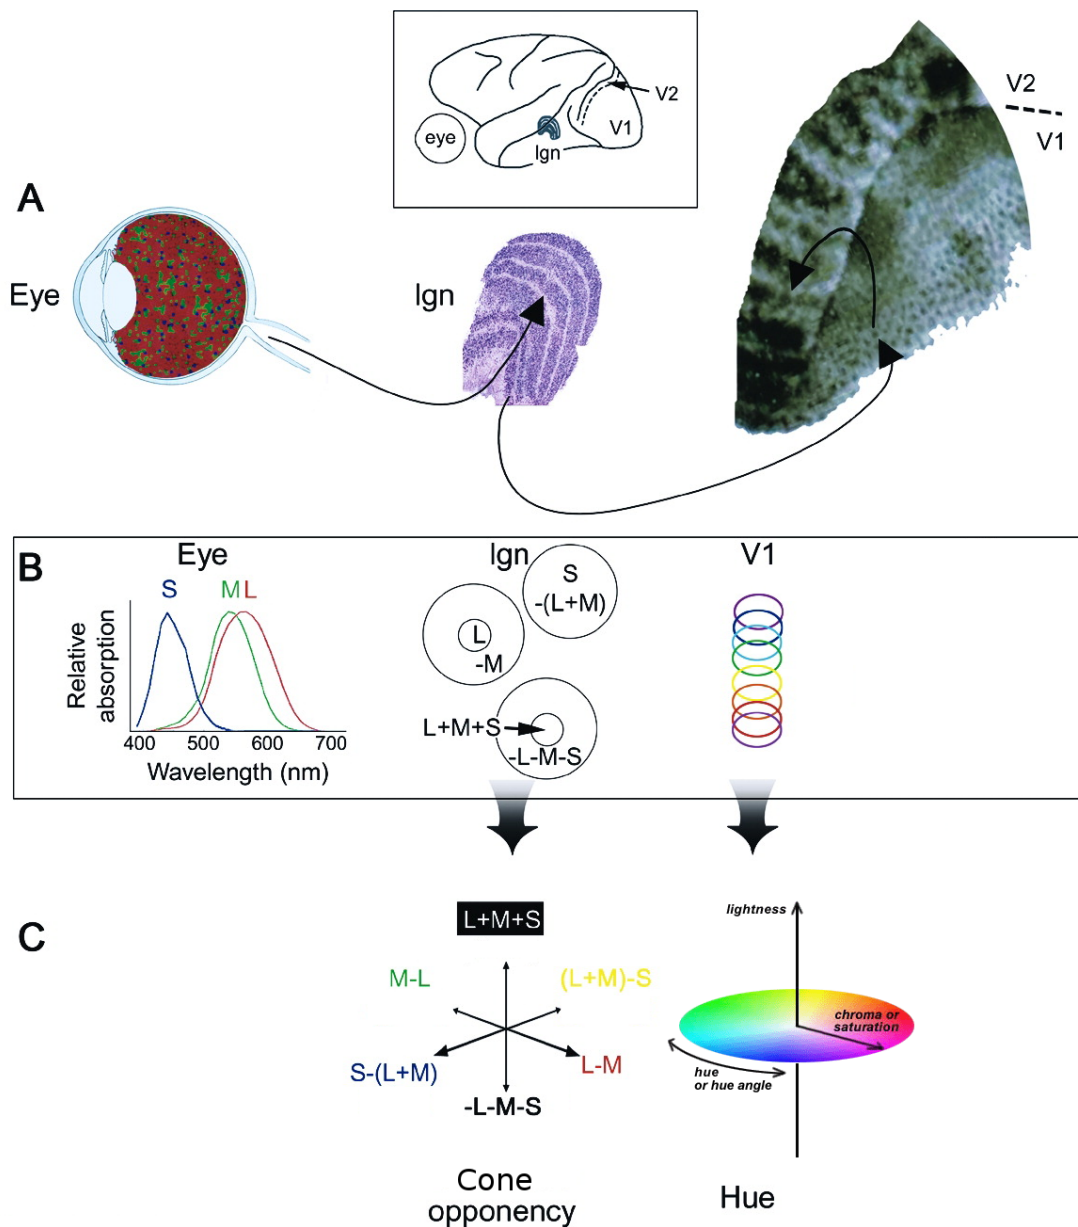
Figure 1.9: *Color representation in the early visual system of primates. Summary of A: anatomy; B: physiology; C: perception. Light is absorbed by three types of photoreceptor, randomly arranged on the retina, and with different but broad and highly overlapping sensitivities. Individual cones cannot distinguish changes in wavelength from changes in physical intensity. Neurons that project from the retina to the lateral geniculate nucleus appear to fall into a few classes, and oppose the output of one cone type against another (e.g. computing an L-M signal). While this allows detection of wavelength changes, these cells do not appear to directly represent perceptual color space. Recent exciting experimental work, however, indicates that there may be spatial organization of cells in V1 according to PCS. Responses to color stimuli are patchy, and within patches there appears to be an organization according to hue. The following chapter will cover this material in more detail. Adapted from Conway (2003).*

organization for PCS:

- Photoreceptors: Three classes of photoreceptor (L, M, and S cones) with different but broad and highly overlapping wavelength sensitivities, randomly arranged on the retina, present in highly variable ratios. The output of individual cones cannot be used to distinguish wavelength from intensity, let alone predict perceived color, and thus the cone activities have no obvious relationship to PCS.

- Retinal ganglion (and lateral geniculate) cells: The output from multiple cones is compared in the cone-opponent retinal ganglion cells, allowing distinction of wavelength from intensity. However, retinal ganglion cells typically cluster around two axes defined by L vs. M and S vs. L+M cone opponency. These cone opponencies describe only limited aspects of our PCS, and there is no evidence for any spatial arrangement based on cells' opponencies.

- V1: Along with the spatially contiguous orientation preference maps introduced above, there appear to be patches of color-responsive cells. Within a color patch, many perceptual hues appear to be represented, and the hues appear to be organized according to PCS.

While recent experiments in macaque indicate that V1 has some organization by PCS, this has only recently been established, and there are still many aspects of this organization that need further testing. V1 is a relatively early and well understood visual area, so results about hue organization in V1 are exciting because they may allow us to link our currently abstract ideas of PCS with concrete and scientifically rigorous neural anatomy and physiology.

1.3.2 How does perceptual color space arise?

In addition to the question of *where* PCS arises, we also do not know *how* the necessary circuitry could arise. Complete genetic specification of PCS seems unlikely: we know that both physiological and environmental variabilities are critical to the development of color vision. For instance, in normal individuals, the region of V1 associated with the foveola (a central region of the retina with no rod photoreceptors, only cones) does not respond to dim light, but in individuals whose cones do not provide any signal, that region of V1 does respond to rods (Baseler et al., 2002). This shows even fairly major physical abnormalities can be adapted to during development. As for the importance of visual input in shaping development, an example is that macaques reared under abnormal color lighting make abnormal color similarity judgments vs. normally reared

macaques (and humans), and completely lack color constancy (Sugita, 2004). Young human infants are unable to perform laboratory tasks dependent on color vision, even though their cones are fully developed (Hamer et al., 1982), suggesting an opportunity for color vision to develop with visual input in humans too.

Intriguingly, however, the color vision system also seems relatively insensitive to both physiological and environmental variability. The ratio and arrangement of the three cone types and their effective sensitivities are variable between color-normal observers (Hofer et al., 2005a; Stockman and Sharpe, 1999), but, despite this, different individuals have broadly similar perceptual color spaces (Webster et al., 2000). Individuals from environments with different color statistics have broadly similar color spaces as well (Webster et al., 2002). Additionally, physiological responses can change over an individual's lifetime—the lens yellows with age, for instance, drastically reducing short-wavelength light reaching the retina—yet color perception remains largely unchanged (chapter 2). This suggests that while the development of the color visual system is driven by visual input, and can be affected by physiological variation, it also adapts to many irrelevant variabilities in physiology and experience. In fact, we know that even the adult color visual system is highly adaptable both to changes in physiology (even incorporating completely new types of color receptor—Mancuso et al., 2009) and to changes in environment (e.g. wearing red goggles for a week greatly shifts color judgments, and leaves color judgments shifted afterward for a similarly extended period until returning to normal; Neitz et al., 2002).

Since it seems that PCS has begun to appear by V1, we wish to know what transformations of visual input by or before V1 could account for PCS, how the necessary wiring could develop based on input and known physiology, and how the system could develop a reasonably stable PCS despite variabilities in physiology and input. Computational modeling of a working visual system can help integrate findings about natural images, the physiology of the early visual system, and human behavior. Can we use such modeling as a tool to show how the wiring necessary for PCS could arise through development, while showing the observed relationship to the physiology and the environment? To answer this, we will build a simple model incorporating known aspects of visual system physiology, and using realistic natural input, to see if the model is able to develop a realistic PCS. Since it is unclear where the necessary early visual system structure comes from in animals, we can generate hypotheses from the model. Along the way, we will see that understanding the neural representation of color is an excellent model for understanding how the brain represents and processes information in general because of the clear and qualitative differences between physical stimuli (wavelengths) and their representation in PCS, as illustrated above. If we can find a

specific location in the brain where the representation corresponds to the perceptual experience, we can understand that circuitry as a model for other aspects of sensory processing.

1.4   Goals of this thesis

The first goal of this thesis is to show that understanding the neural representation of color is a useful model for understanding how the brain represents and processes information in general (*goal 1*). The second goal is to integrate findings about natural images and the physiology of the early visual system, in the form of a simple computational model, to see whether such a model can develop a realistic perceptual color space (*goal 2*). By analyzing the model, whether it successfully develops a realistic perceptual color space or not, we should be able to generate hypotheses about where the necessary early visual processing comes from in animals (*goal 3*). The final goal is that the work done in this thesis should suggest future experimental work, and lead into future modeling work, which together can be used to evaluate the hypotheses generated (*goal 4*).

To achieve these goals, we demonstrate from existing work that in humans and macaque a reliable perceptual color space arises from highly *un*-reliable underlying physiology, and an environment that does not reflect perceptual color space (chapters 2 and 3). In chapters 4 and 5, we identify five forms of environmental and physiological variabilities not accounted for by current models, preventing the models from developing a realistic representation of PCS, and we create a new model that *does* develop a realistic representation of PCS. Below we outline our work in more detail, chapter by chapter.

1.5   Chapter summaries

- Chapter 2 outlines the background necessary for this thesis, and introduces important themes that our subsequent modeling work will tackle. We give an outline of color vision, examining the perceptual experience of color, input to the color visual system, and the biological processes in between the two that perform the transformation. We review how visual input and physiology have a critical impact on the development of color vision, but at the same time we emphasize large variabilities in both the physiology and environmental input that seem to have a surprisingly small effect on the development of perceptual color space.

In more detail, we review experimental results about orientation and color representation in macaque V1, because this is the first area of the early visual system in which there appears to be any organization according to PCS. Finally, we review previous studies that have modeled the development of V1's representation of orientation and color. We show how these models have been successful in explaining some aspects of the data, but that importantly the models do not explain key experimental findings about the spatial organization of hue in V1.

- The previous chapter showed that visual input is critical to the development of color vision in animals. In chapter 3, we first review previous investigations of the color statistics of natural scenes, in order to compare the distribution of color percepts experienced with perceptual color space. We conclude that colors in natural scenes appear to be highly unrepresentative of PCS. Since in subsequent chapters we will be modeling the input-driven development of PCS in V1, we need both to simulate the visual experience of an individual, and to understand the color statistics of this visual input. We take natural scenes to be representative of visual experience, and analyze the color statistics of a number of databases of calibrated color images of natural scenes. In agreement with previous analyses of natural scenes, we find that these databases contain a highly non-uniform sampling of perceptual color space, with some colors being present in large quantities, while others rarely or never occur. Experimental evidence from animals is for a relatively perceptually uniform representation in V1 and secondary visual cortex, so this raises an interesting problem for input-driven developmental models. How can such a model learn a reasonably uniform perceptual color space from such biased input?

- Having seen in previous chapters that it is not obvious how an organization of color cells by perceptual space should arise (WS and PCS are quite different; visual experience is a highly non-uniform sampling of PCS; there is a large amount of physiological variety), in chapter 4 we begin to model the development of the early color visual system. We start from the modeling work reviewed in chapter 2, first addressing a number of problems that prevent previous models from developing orientation and color maps in a realistic fashion. The result is our first model, which we call the *naive model* , and which—having improved the previous most successful model—becomes the best available model of the development of orientation and color selectivity in V1. However, we show that while the model develops a realistic spatially contiguous orientation map and forms color (CR) patches, as in macaque, the organization of hue representation does not agree with data from macaque. Crucially, the naive model does not

support coding of hue based on spatial location. We speculate that the model, unlike the visual system, is not able to adapt to important variabilities in input and physiology.

- The challenge in chapter 5 is to create a model like the naive model, but that *can* replicate the experimental hue organization results. We do this by analyzing the naive model to determine what prevents a realistic hue organization arising, finding it is indeed sources of variability and bias that the naive model has no mechanisms to overcome—unlike the actual visual system. We show by constructing a model without these variabilities and biases—the *idealized model* — that it is possible for a computational model to develop a realistic organization for orientation and color, including a realistic hue representation. In creating this model, we demonstrate important problems that the brain appears to be overcoming in constructing PCS, but that previous models are not, such as the input being a highly non-uniform sampling of PCS, and variable ratios of cone photoreceptor types. We speculate that well-accepted biological mechanisms such as neural adaptation are what allows the brain to be surprisingly insensitive to these problems. These mechanisms have typically been omitted from models, but this chapter demonstrates that, at least for color, they are critical and cannot be omitted.

- Chapter 6 discusses implications of the above findings, including suggesting new experimental work, and proposing experiments to perform on the previous chapter's models, as well as the next steps for extending the modeling work.

- Chapter 7 summarizes the thesis, and evaluates its contribution.

Chapter 2

Background

2.1 Introduction

This chapter begins by giving an overview of the steps required for color vision to happen, starting from light leaving a source, through measuring its representation in the brain, and finally to behavioral measurements of human perceptual color space. Along this pathway we identify key stages to describe in more detail in subsequent sections:

- Our experience of color: describing color, and our perceptual color space.

- Input to the color vision system: wavelength space.

- The biology transforming wavelength space to perceptual color space: the early visual system in between. These aspects are less well established than the two above.

While reviewing the early visual system, we find that primary visual cortex (V1), alongside its well-known role in orientation (OR) coding, appears to be the first area of the brain representing color (CR) in a way related to perceptual color space (PCS). Following this, we focus in detail on experimental results about the representation of OR and CR in macaque V1. Finally, we review existing models of OR and CR organization in V1 to see how much of the organization of V1 has been explained computationally.

2.1.1   Contributions of this chapter

- A review of the early visual system for color vision, emphasizing the development of a stable PCS despite large underlying physiological and environmental variabilities.

- Comparison of different experimental investigations of CR and OR representation at the V1 map level, demonstrating that results are broadly consistent in showing CR cells organized into isolated patches within a spatially contiguous map of OR (figure 2.20), and an organization of hue within CR patches that follows PCS (figure 2.23).

- Review of previous computational models of OR and CR development in V1, showing they can explain the organization of OR maps and CR patches, but cannot account for data showing an organization for hue according to PCS.

2.2   Overview of color vision

Here we give an overview of the steps that take place to enable color vision, following figure 2.1. The story is necessarily simplified; we focus on steps relating to the key experimental and behavioral data that we will present later. We begin with a light source, such as the sun. Light can take several different paths to the eye, but two typical, simple paths occurring when we are outside in daylight are:

1. From the source, light is transmitted through air and then the eye's media (including cornea, lens, vitreous humor, aqueous humor, blood vessels, macular filter, and even nerve cells) before reaching the light-sensing cells (photoreceptors).

2. Light, which has come from a source, is reflected by an object's surface and is then transmitted, as above, through air and the eye's media before reaching the photoreceptors.

We are ignoring many other interactions of light and matter, including interference (e.g. brown-pigmented peacock feathers appearing blue or green), diffraction (e.g. the many colors that can be seen in "the play of color" from opal), atmospheric scattering (e.g. a midday blue sky vs. a "golden hour" red sky), dispersion (e.g. a rainbow), transmission through solid objects (e.g. through leaves in a forest, or tinted glass), and more (see Nassau 2003 for an overview).
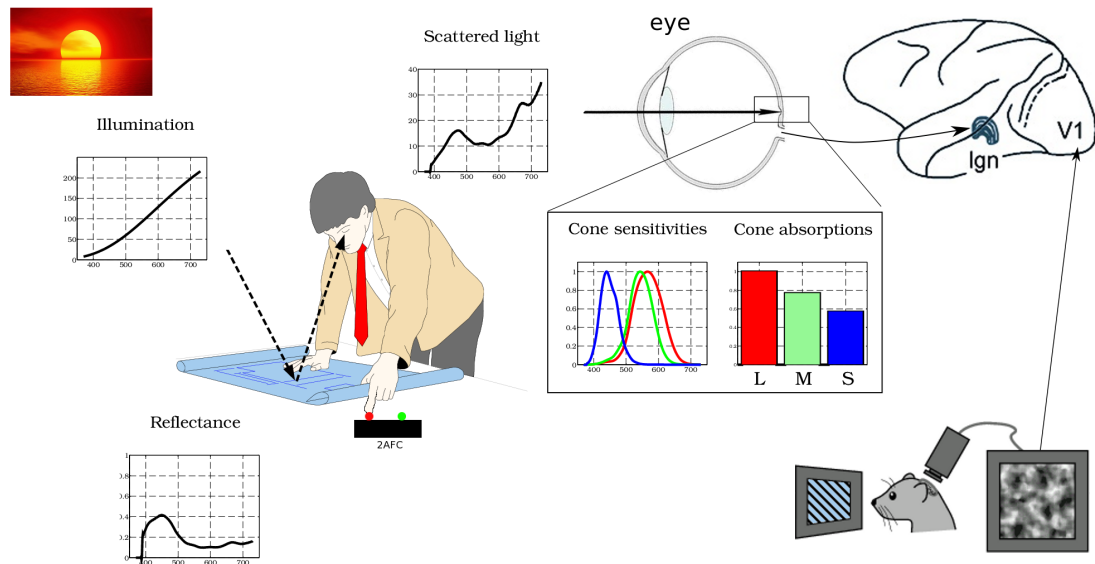
Figure 2.1: *Color vision pathway. Cartoon overview of the steps described in the text. Illumination from a source is reflected by an object, and the scattered light is observed by an individual. The individual could be judging whether the object is the same color as one viewed just previously in a 2AFC test, for instance, allowing us to learn about perceptual color space. Meanwhile, the process occurring in the early visual system is that light is absorbed with some probability by one of three cone types. The different cone types, L, M, and S, have different but broad and highly overlapping spectral sensitivities. The signals from the cones pass to the retinal ganglion cells, which project (via the optic nerve) to the LGN. From there, the signal passes to V1. Activity in the brain may be recorded with various techniques, some invasive and others not. Here we show optical imaging of animal V1, which is the method used to gather most of the results about OR and CR organization that form the basis of this thesis. Note that the picture here is simplified in many ways, including that spatial information is not shown (e.g. the scattered light will vary dramatically by location except in deliberate laboratory conditions). Adapted from Wandell (1995) and Miikkulainen et al. (2005).*

For daylight color vision—photopic vision, the focus of this thesis—the photoreceptors are the three classes of cone on human (or macaque) retina. The three cone types have different but broad and highly overlapping spectral sensitivities. Once light has finally reached photoreceptors, it may be absorbed by a photoreceptor, causing isomerization of a photopigment and the start of a chain of processing. This chain of processing leads out of the retina via the retinal ganglion cells, which project (via the optic nerve) to the LGN, near the center of the brain. In the simple, feed-forward view of color processing we discuss here, the LGN does not significantly change the color representation from that of the retinal ganglion cells. The LGN then projects mainly to V1 (at the back of the brain, on the cortical surface), a thin ($2\,\mathrm{mm}$ thick vs. $>90\,000\,\mathrm{mm}^2$ area), folded, layered sheet of neurons and connections. Input arrives in layer 4 (L4), after which, "output" to higher visual areas proceeds through higher layers (with lower layer numbers), and "feedback" to lower visual areas proceeds through deeper layers.

As shown in figure 2.1, our pathway splits at V1: the imaging results we will study in a later section view supragranular layers (L1–3), 0–400 µm below the surface. After V1, the visual pathway continues in the brain through more areas, and somewhere our conscious perception of color arises—but this is not something we know how to measure directly. However, we do have experiments that use behavioral outputs to give information about our perception, e.g. 2AFC experiments judging whether or not two colors are different, which can help to construct a perceptual color space through measuring just-noticeable differences (jnds).

Having seen a simplified overview of our color vision pathway, we now describe in more detail the end of the pathway most familiar to us: PCS. Afterward, we will go on to describe the start (incoming light, wavelength space), before focusing on the less well-established biological pathway in between (the early visual system).

2.3   Describing color

There are many ways to describe color. For instance, the color of a monochromatic light against a black background can be described by its wavelength. However, in the previous chapter we saw that not all our color sensations can be recreated with a single wavelength. Additionally, even the color of a monochromatic light can depend on luminance (Purdy, 1931), as well as on the observer's state of adaptation (to be discussed in section 2.4.3). The color of light could instead be described by the L, M, and S cone excitations it causes—a receptor space (figure 2.2). However, we will see in
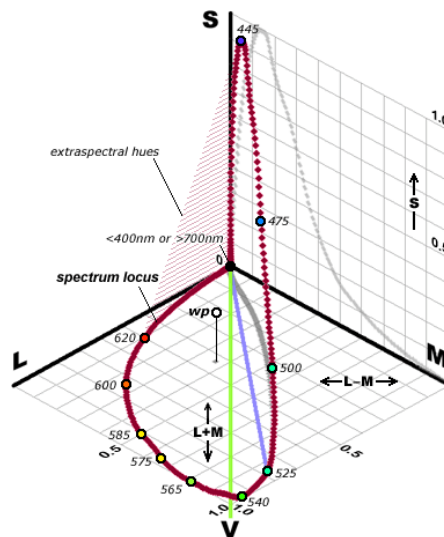
Figure 2.2: *Color spaces 1: cone excitations. Monochromatic light of variable wavelength plotted as cone excitations, i.e. light described by the excitations it causes in the three cone types. The shape traced out by such a light, the "spectrum locus", is determined by the overlapping cone spectral sensitivities (figure 2.1, and shown in more detail in a later section). For instance, short-wavelength (e.g. $445\,\mathrm{nm}$) light strongly stimulates S cones, but has little impact on L and M cones, while on the other hand, the M cone cannot be stimulated in isolation. However, this representation has a number of shortcomings. How should we scale the axes? Some humans have far more L than M cones, while for others it is the other way round. Also, as described later in the text, the same stimulus can cause different LMS responses (depending on context), and the same LMS excitations can appear different colors. Note that a range of colors we perceive, the "extraspectral hues", cannot be stimulated by a single monochromatic light. Image reprinted from MacEvoy (2009).*
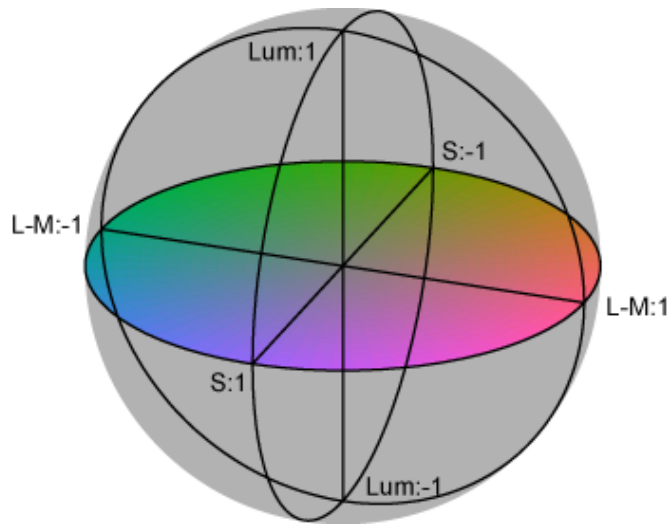
Figure 2.3: *Color spaces 2: cone opponency. The color space shown is DKL (Derrington et al., 1984); note that the colors shown in this printed representation are only an approximate guide. Coordinates in this space represent the responses of the three assumed retinal ganglion cell classes; the axes are often referred to as the "cardinal axes".*

section 2.5.1.2 that different perceived colors can cause the same receptor excitations, and, conversely, different receptor excitations can appear the same color, so responses in such receptor spaces do not adequately describe our perception of color. A third possibility for describing color is to use a cone-opponent space (figure 2.3), reflecting the typically assumed cone opponency in the RGC and LGN (e.g. using L vs. M and S vs. (L+M) axes, which will be described more fully in section 2.5.2). These two axes capture only a limited aspect of our perceptual experience, though.

Before the cone spectral sensitivities were known, it was found that the perceived color of any light viewed in isolation could be matched by a variable mixture of three monochromatic, "primary" lights (i.e. with different, independent wavelengths; figure 2.4). The CIE 1931 XYZ color space (CIEXYZ) color space is based on results of observers adjusting mixtures of three primaries until they match test stimuli. The CIEXYZ color space thus allows us to tell whether two mixtures of light will look the same (in the same viewing conditions), but does not tell us how they will actually appear. Figure 2.5 shows the *(x,y)* diagram, which is constructed from CIEXYZ by normalizing out luminance information.

All the color spaces described so far are important because they are in widespread use, and are unavoidable in the color and physiological literature. However, we can find a better description of our perceptual experience of color in spaces such as the Munsell color system (technically a color order system), CIE 1976 L* a* b* color space
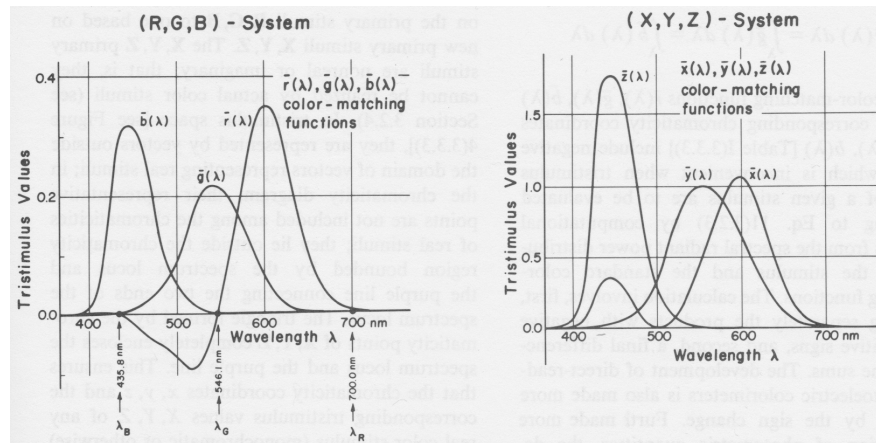
Figure 2.4: *Color spaces 3: CIE 1931 colorimetric system. Before the human cone photoreceptor sensitivities could be measured, it was at least known that there were three different receptor systems involved in human color perception (trichromatic theory), and that the color of any particular light could be matched for a color-normal human observer by mixing three independent light sources in various proportions (a consequence of Grassmann's laws; Krantz, 1975). By independent, here we mean that no combination of any two can match the third.* Left-hand graph*: indicates the amounts of three particular monochromatic lights ("primaries") that are required to match the wavelength of a test light. Note that negative values of a particular primary indicate it must be added to the test light in order for the observer to make a match, and is a consequence of the primaries not being able to produce a saturated enough color. Based on the average of results from a number of color-normal observers, the $\bar{r}(\lambda)$, $\bar{g}(\lambda)$, and $\bar{b}(\lambda)$ functions allow the specification of color in terms of three numbers (tristimulus values), for any incoming spectral power distribution $P(\lambda)$ (e.g. the first value would be $R \propto \int \bar{r}(\lambda)P(\lambda)d\lambda$). Any lights that produce the same tristimulus values will appear identical for an observer (under the same viewing/observer adaptation conditions).* Right-hand graph*: the standard CIE color system is not based on real primaries, but on imaginary primaries that would be able to match all possible colors; i.e. primaries that are more saturated than it is possible to create with monochromatic primaries (ultimately because, for instance, it is not possible to stimulate M cones without stimulating L or S cones). The standard $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$ functions, a linear transform of the $\bar{r}(\lambda)$, $\bar{g}(\lambda)$, and $\bar{b}(\lambda)$ functions, therefore have no negative components. The resulting XYZ tristimulus values are the current standard in color science for communicating device-independent color information. Graphs reprinted from Wyszecki and Stiles (1982). Note that a detailed re-derivation of the XYZ system from the original RGB matching data is available in Fairman et al. (1997).*
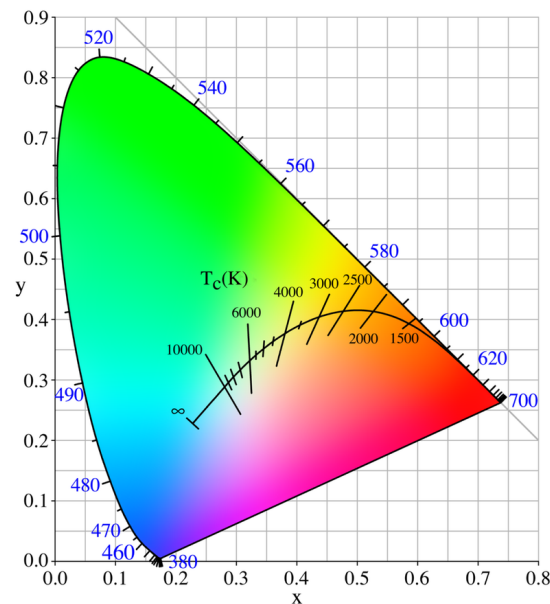
Figure 2.5: *Color spaces 4: CIE 1931 (x, y) chromaticity diagram. The* (x,y)
*diagram is important and ubiquitous in color literature. Viewed in isolation,
the color of a monochromatic light, or mixtures of spatially homogeneous
monochromatic lights, can be specified by three tristimulus values (not the
L, M, S values; see figure 2.4), and normalizing away the luminance infor-
mation $(X + Y + Z)$ allows two coordinates (x and y) to be used to specify
chromaticity. The diagram helps to show which mixtures will look the same,
although not how they will actually look (and hence the colors shown in the
diagram can be misleading and are there only as an introductory guide). As
examples of reading the diagram, various types of daylight are marked by the
black line. However, the* (x,y) *diagram is far from uniform; equal distances
in different parts of the diagram have very different perceptual differences
(compare the area appearing green to that appearing yellow). Further, for
specifying color appearance, the chromaticity coordinates are problematic
because the observer's adaptation to the illuminant is not taken into account,
and hence the color percept from a chromaticity coordinate can vary. We will
avoid using chromaticity diagrams to describe color appearance, except to
discuss historical data presented on such diagrams.* Image "Planckian Locus" by
PAR (via Wikimedia Commons) is in the public domain.

(a) Munsell   (b) CIELCh$_{ab}$ (c) HSV
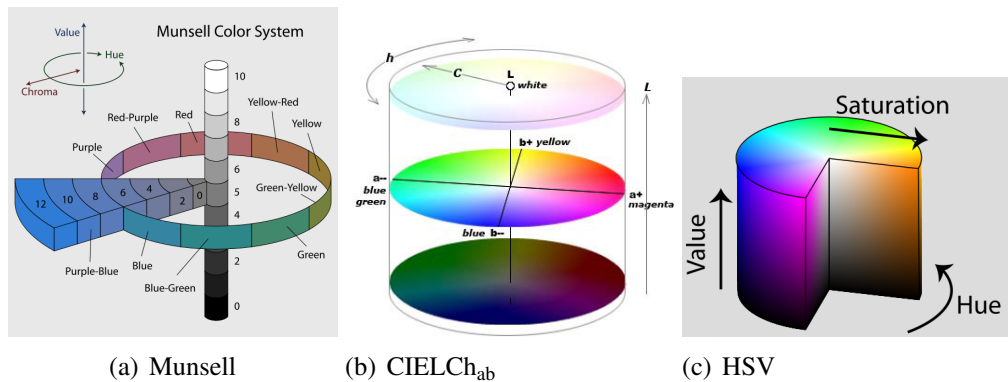
Figure 2.6: *Color spaces 5: perceptual. These spaces describe three perceptual aspects of color: hue (red, yellow, green, etc.), colorfulness/chroma/saturation (not all the same quantity, but approximately how intense/different from gray the color is), and brightness/lightness (again, these are not all the same quantity, but approximately perceived total amount of light). Hue is circular, in that if people judge how similar different mixtures of monochromatic lights are to each other, the result will be a circle. Lightness and chroma are not circular. As discussed in the text, perceptual color space is in fact more complex than shown here; additional variables are required to fully describe color. Also note that because of limitations of color printing and display devices, the colors shown may not actually be perceptually uniform, and as in earlier figures are intended only as a guide. (a) Image "The Munsell color system" by J. Rus (via Wikimedia Commons) is licensed under CC BY-SA 3.0. (b) Reprinted from MacEvoy (2005). (c) Crop from image "Hsl-hsv models" by J. Rus (via Wikimedia Commons) is licensed under CC BY-SA 3.0.*

(CIELAB) and CIE 1976 L* u* v* color space (CIELUV) (uniform color spaces), or even HSV (from computer graphics), which all have the concepts of three separate perceptual attributes of color: hue, colorfulness (or chroma), and lightness (or brightness)—see figure 2.6. Some of these perceptual color spaces have been constructed to be perceptually uniform; that is, equal distances in the color space should correspond to equal perceptual differences. Such a space could possibly be built from human experiments, e.g. using jnds (Takamura and Kobayashi, 2001) or judgments of similarity (Shepard (1962), as cited in Brainard (2003)—but see Wuerger et al. (1995) for problems with the multidimensional scaling method employed). In practice, though, spaces in use now are usually derived from previously defined spaces. For instance, CIELAB (figure 2.6(b)) is a nonlinear distortion of an opponent space based on CIEXYZ, and includes a simple model of adaptation to viewing conditions. The transformations were developed so that color differences slightly larger than jnds would be the same in all parts of the hue circle. This means e.g. a change in hue angle of e.g. 5° in the "red" part of the hue circle ought to have the same perceptual impact as a 5° change in the "green" part of the hue circle. Additionally, a change to one dimension (e.g. lightness) should not affect other dimensions (e.g. hue). While in CIELAB these are certainly closer to being true than in e.g. HSV, true uniformity was not completely achieved. Large differences in color, for instance, are distorted (Fairchild, 2005). In fact, there are several "uniform color spaces", because no space so far has achieved these aims, and it is clear that we do not yet fully understand the geometry of color space (Zaidi et al., 2013).

Apart from not being truly uniform even in the restricted viewing situations to which they are intended to apply, spaces such as CIELAB are too simple to be true color appearance models. For example, CIELAB and similar spaces do not deal with related colors (figure 2.7), or with the effects on color of shape and size. Extensions of CIELAB exist, such as sCIELAB (Zhang and Wandell, 1997), which includes spatial filtering in LMS (i.e. considers a pixel surrounded by other pixels). There is also CIECAM (Fairchild, 2005), which models color appearance to predict lightness, brightness, chroma, hue, saturation, and colorfulness. However, in this thesis we will go no further than using and modeling the representation of three-dimensional spaces such as CIELAB and CIELUV, or their two-dimensional (hue-only) representations without luminance (figure 2.8).

Even relatively simple perceptual color spaces like CIELAB cannot currently be linked back in any obvious way to underlying physiology (whereas the earlier physiological spaces cannot easily be related to color perception). This is one of the reasons for excitement at apparently finding an organization according to perceptual hue in V1 (in-
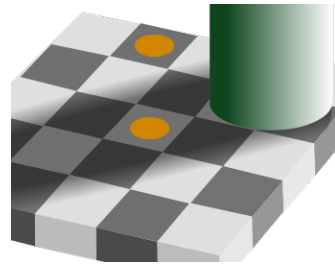
Figure 2.7: *Related colors. To appear brown instead of orange (or vice versa), a surface (or light) must be appropriately surrounded (either in space, as shown, or in time). Both disks are physically the same, but one appears orange and the other brown. There is no wavelength of light that is "brown" in isolation, rather than "orange".* Image "Optical illusion" by JunCTionS (via Wikimedia Commons) is in the public domain.

troduced in the previous chapter, and discussed further in a later section). Compared to higher cortical areas, V1 is relatively easily accessible, and relatively well understood (see e.g. Olshausen and Field, 2005, or Carandini et al., 2005). This gives the possibility of linking our currently abstract ideas of PCS with concrete neural anatomy.

As first stated in chapter 1, one of the fundamental questions we aim to contribute towards answering is how incoming light is transformed into a perceptual color space. We have seen that PCS is quite different from wavelength space (WS), and that it is not obvious how to transform between them, which is already an interesting question. However, there is another puzzle relating to the development of PCS in the brain, also introduced in chapter 1: PCS is broadly similar between individuals, despite early visual system physiology and visual experience varying between individuals. PCS is also reasonably constant for one individual over time, despite physiology and visual experiences that change over time. This constancy exists despite the fact that early visual system physiology and visual experience both have a critical impact on the development of PCS. In the following sections, we will review input to the visual system, and the biology of the early visual system; throughout, we will report results that show the impact of visual input and physiology on the development of color vision and PCS. In some cases, the impact is large, while in others it is small—the mixture is surprising.

2.4 Input to color visual system

In our overview of the color vision pathway, we saw that incoming light is somehow transformed to PCS in the brain, and in the previous section we looked at PCS in more detail. In this section, we look at the input in more detail. The perceptible input
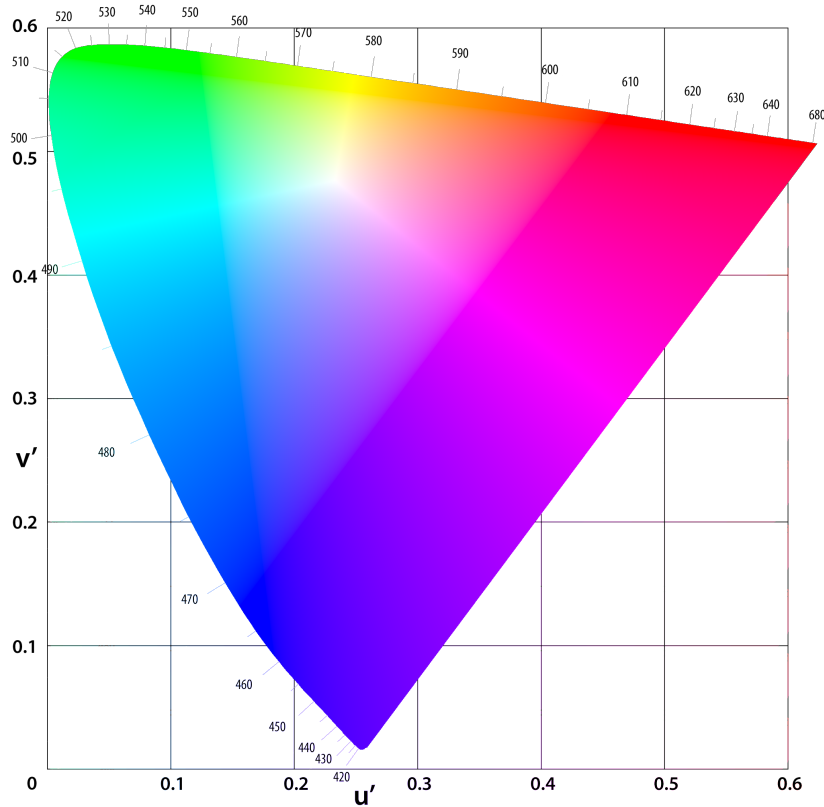
Figure 2.8: *Color spaces 5: uniform chromaticity diagram. Earlier, in 2.5, we saw a two-dimensional constant-luminance slice through XYZ: the (x,y) diagram. Some data in later chapters will be presented on a similar diagrams, but based on improved (more perceptually uniform) color spaces (e.g. LAB and LUV; figure 2.6). Shown here is the UCS diagram: luminance has been normalized out from LUV, and a two-dimensional chromaticity diagram produced (like figure 2.5, but intended to be more percetually uniform). Note that because of limitations of color printing and display devices, the colors shown may not actually be perceptually uniform, and are again intended only as a guide. Image "CIE 1976 UCS" by Adoniscik (via Wikimedia Commons) is in the public domain.*

is a tiny fraction of a wide range of available electromagnetic radiation, specifically that between about 400 and 700 nm (the wavelength range to which the photoreceptors are sensitive, after filtering by the eye's media). At any point of the eye, a spectral power distribution (SPD), varying over time, is arriving. This incoming signal, WS, is quite different from our perception of color. For instance, a color-normal individual viewing monochromatic lights in isolation would likely judge 400 and 700 nm light to be more similar to each other than either is to 550 nm light. We listed a number of other differences between WS and PCS in chapter 1. Despite this difference between WS and PCS, later in this section we will argue that incoming light plays a key role in the development of color vision. At the same time, however, the visual system seems surprisingly unaffected by certain sources of variation in the incoming light. This motivates a desire to understand the properties of the incoming light, characterized as long-term statistics, in order to be able to model how PCS could arise in a model whose development is driven by visual input.

To characterize the statistics of incoming light, there are a number of aspects we could consider. For instance, what distribution of energies does light have relative to the LMS cone sensitivities—are some cone types stimulated more, on average, than others? What about the combined spatial and spectral information in the incoming signal— how does the spectral content vary over space? Or, we can consider the colors we perceive: do we experience some hue percepts more than others in natural scenes? To understand the development of the color visual system, it is important to understand various properties of the input. Additionally, to model this development, we will need a concrete way to simulate the visual experience of an individual. Therefore, the next chapter goes into detail on this topic, before we later go on to modeling in subsequent chapters. Here, we return to the impact of visual input on the development of the color visual system.

2.4.1 Environment can greatly affect development of color vision

Incoming light plays a crucial role in shaping the visual system. For instance, deprivation of input to one eye of a kitten during a critical period of development leads to the kitten later being blind to input from that eye (Wiesel and Hubel, 1963). Another example is that restricting the orientations experienced during a kitten's development strongly affects its ability to perceive other orientations after the restriction is removed (Blakemore and Cooper, 1970). Similarly, deprivation of motion leads to no direction selectivity (Cynader and Chernenko, 1976; Li et al., 2006). In all these cases, the

eye and LGN still appear normal—it appears to be the visual cortex that is affected. Similarly, human visual development appears to be affected by visual input, too. For instance, a human may develop a stereoanomaly or stereoblindness if during development normal binocular vision is not possible (e.g. because of strabismus or amblyopia; Fawcett et al., 2005). Note that while there may be "critical periods" during which development is particularly sensitive, it appears that many aspects of vision can be similarly affected in adults, although perhaps more slowly or to a lesser degree (to be discussed in section 2.4.3). Additionally, many of the effects seen during "critical periods" involve changing visual circuitry already established at birth, not necessarily preventing its initial development (reviewed in Huberman et al., 2008).

Visual input has a major effect on the development of color vision too, although this has been difficult to show (Brenner and Cornelissen, 2005; Wagner and Kröger, 2005), at least partly because color vision has been less well understood (for instance, the difference between wavelength discrimination vs. color vision as discussed in the introduction). However, there are examples of key aspects of the physiology underlying color vision being modified by visual experience. For instance, McCourt and Jacobs (1983) showed that the wiring of squirrel retinal ganglion cells is modified postnatally, and varies depending on rearing condition (red light, white light, or darkness). When blue acara fish are reared with only short-wavelength light for input, their short-wavelength-sensitive cone population decreases, and long- and medium-wavelength-sensitive cone outer segments lengthen (Kröger et al., 1999; Wagner and Kröger, 2000). The responses of horizontal cells (which connect photoreceptors) also change (Kröger et al., 2001). These changes were accompanied by behaviorally detectable changes in color vision.

What about the effect of visual experience on the development of primate color vision? Human infants over two months are able to perform a variety of laboratory tasks dependent on wavelength discrimination, whereas infants under two months are not able to perform these tasks, despite having functional photoreceptors (Hamer et al., 1982; Teller, 1998). For instance, blind-to-the-stimulus observers can successfully identify which half of a color screen seven-week-old infants are preferentially looking at in response to a color test pattern presented randomly on either half of the screen (test and background matched for each infant's luminous efficiency by requiring success across multiple radiances), but are unable to do so with three-week-old infants (Clavadetscher et al., 1988). While such results suggest visual input has the opportunity to shape human PCS, stronger evidence of visual experience's impact on the development of color vision comes from animal experiments. For instance, it has been possible to show that color vision development can be modified significantly by experience in macaque

monkeys. Sugita (2004) raised macaque monkeys from one month until one year old in a room with special lighting. For 12 hours a day, the room was illuminated with monochromatic light whose wavelength changed once per minute, randomly switching between one of 465, 517, 592, and 641 nm. After one year, the monkeys were moved to a room lit by daylight. These selectively reared macaques—alongside normally reared macaques, and humans—were then tested on their ability to make brightness similarity judgments, color similarity judgments, and their color constancy was evaluated. Color similarity judgments for the normally reared macaques (and for the humans) showed the typical result that colors were grouped into reddish, greenish, and blueish categories (humans and normally reared monkeys rarely judged red, green, or blue to be similar). The selectively reared macaques, however, made very different similarity judgments, often judging red, green, and blue as similar. This suggests that the selectively reared monkeys either did not develop the PCS that humans and normally reared macaques appear to have, or such organization was overwritten by the rearing scheme.

Neither the selectively nor normally reared monkeys had trouble performing a brightness similarity matching task on achromatic stimuli, but while the normally reared monkeys could immediately transfer successfully to the same task on chromatic stimuli, the selectively reared monkeys required 30 days of intensive training (while living under daylight) to perform at the same level.

Finally, the selectively reared monkeys had no color constancy. Both groups of monkeys (and human observers) could identify a red target rather than a yellow target against a Mondrian background (multiple sizes and colors of rectangles) under a white light, but when the light was changed such that the yellow target would appear red in isolation under the light, the selectively reared monkeys could no longer identify the red target (they selected the yellow target) while the normally reared monkeys (and human observers) could. Even after nine months of training at this task, the selectively reared monkeys did not improve.

2.4.2 Environment has a surprisingly small impact on development of perceptual color space

Having seen that the development of the color visual system depends critically on the environment, we now come to the surprising finding that variabilities and biases in the environment paradoxically seem to have only a relatively small impact on the development of PCS.

Given the evidence outlined above that color vision is affected by visual experience, one might expect that our perceptual color space would reflect our experience with colored stimuli in the environment. However, in the next chapter, we will review previous work indicating that natural scenes contain a very limited range of PCS, and represent a highly non-uniform sampling of PCS. We also analyze the distribution of colors present in databases of calibrated color images of natural scenes, confirming this finding. (When we talk about the colors of scenes, we mean the percepts generated in a human in response to those scenes, using a model of adult color perception based on calibrated image data.) Somehow, despite the large difference between PCS and visual experience, and visual experience having a major impact on the development of color vision, the brain still develops PCS.

Even more surprisingly given these variabilities, humans from different locations and cultures seem to share a broadly similar perceptual color space (Berlin and Kay, 1969), and the perceptual color spaces we described earlier work well practically. However, the details of color space do appear to vary between individuals. A standard way in which individual color space has been investigated is by testing "unique hues" for an observer; that is, hues which appear "pure". Thus, for example, "unique yellow" is the hue which appears "neither reddish nor greenish" (Hurvich and Jameson, 1957). This method has been supported more recently by non-verbal methods (Logvinenko, 2012), but the majority of data involves unique hues as defined above (and typically with lights rather than surfaces). For some unique hues, the wavelengths chosen by different observers are found to vary more than for others. For instance, unique yellow can differ by 20 nm while unique green by 80 nm (Webster et al., 2002), although the ranges found by different studies vary considerably (Kuehni, 2004, 2014). When the unique hues of groups of individuals from different environments were tested, the unique hues were found to differ significantly between groups, though the differences between groups were smaller than the differences within one group (Webster et al., 2002). However, it is not clear how the environmental statistics differed, or how much cultural factors contributed.

Therefore, despite the environment having a large effect on the development of color vision in some circumstances as we saw earlier, PCS does appear to develop reasonably consistently, but with variability. We will see later when we review the biology of the early visual system, that PCS additionally develops reasonably consistently despite large physiological variabilities. Before that, in the next section, we will see that one explanation for the relative consistency the visual system is able to achieve is likely to be adaptation.

2.4.3 Chromatic adaptation

Adaptation is an important feature of the early visual system, and of the brain in general, in which sensitivity is adjusted to match the available output range (e.g. a limited range of neural firing) to the input domain (Wark et al., 2007). As an example, the 10^4 difference in intensity levels between highlights and shadows that can occur in a visual scene exceeds the range the neurons can signal by a factor of 10^2 (Dunn et al., 2007; Rieke and Rudd, 2009). Multiple processes underlie this particular example of adaptation, and in general, adaptation may involve changes within neurons themselves, or in connections between them, or in the growth of new neural circuitry (Webster, 2011). We will discuss the neural mechanisms potentially underlying chromatic adaptation in later sections; here we will focus on briefly describing findings about chromatic adaptation phenomena relevant to the development of PCS. We will separate adaptation into "short" (seconds, minutes, or hours) and "long" (hours, days, or weeks) timescales, although it is not always clear what constitutes "long" and what "short", and the actual situation is likely to be more nuanced. Also, for now we will focus on adaptation to incoming light, rather than on adaptation of one part of the early visual system to the variable physiology of another (which we will cover later). And as earlier, we are mainly interested in phenomena at daylight levels of vision.

In the previous chapter, figure 1.3 showed some examples of short-term chromatic adaptation. The first example illustrated "discounting the illuminant", i.e. some form of adaptation to the average color in a scene. This is often explained by von Kries normalization (Fairchild, 2005): independent multiplicative scaling in each of the three cone classes based on the mean excitation of each type for a scene, or on the mean of excitation of one type over a smaller area (Chichilnisky and Wandell, 1995; Shevell, 2003), or over recent visual history (Lee et al., 2012). This process of normalizing allows a better representation of contrast, and is an example of defining a cone contrast space. Similarly to how the visual system adapts to luminance contrast as well as to mean luminance, the color visual system also adapts to chromatic contrast, but we will not consider this further in this thesis.

Von Kries normalization is not able to account for all short-term chromatic adaptation effects. Many models of short-term chromatic adaptation additionally include a second stage of processing after von Kries normalization, based on adaptation at post-receptoral, color-opponent processing stage (Jameson et al., 1979). Again, though, we will not consider this further in this thesis because these models are difficult to relate to underlying biology. However, it seems highly likely that adaptation at the RGC and V1 levels is required in some form to account for many observed short-term chromatic

adaptation phenomena (Belmore and Shevell, 2011). The cortex is indicated particu-
larly in chromatic adaptation phenomena involving more spatially complex patterns.

As well as adapting in the short term, the adult color visual system is also highly
adaptable to input over much longer periods. For instance, wearing red goggles for a
week leaves color judgments significantly shifted from their original location. These
shifts remain afterward for a similarly extended period, gradually returning to their
original location (Neitz et al., 2002; Belmore and Shevell, 2008; Eisner and Enoch,
1982). Neitz et al. (2002) included measurements indicating that cone sensitivities did
not change and hence were not the source of the effect. Additionally, adapting only
one eye resulted in the same effect in the other eye, but at a reduced magnitude. This
implies adaptation of cells receiving input from both eyes, implicating the cortex.

Finally, adaptation is not restricted to color in isolation of other features. For instance,
the McCollough effect (McCollough Howard and Webster, 2011) is an aftereffect that
combines orientation and color. This is interesting, because the first cells that show
orientation tuning in humans and macaques are in V1. Additionally, the McCollough
effect is far longer lasting for a given induction period than are purely color aftereffects
(e.g. a ten-minute induction can last for 24 hours).

We now return to our color vision pathway. We started with our perceptual experience
of color (PCS) and then skipped backwards to discuss light coming into the eye (WS).
We saw that WS is not like PCS; the transformation must occur somewhere in the
brain, and in fact possibly early on in the visual system given results from V1, as we
will see in the next section.

2.5   Early visual system

In this section, we will introduce the early visual system, focusing on what is known
about how it could contribute to the development of PCS. Hence, while color vision is
not the only purpose of the mechanisms we will be looking at, we will approach our
review from the point of view of color vision.

A goal of this thesis is to understand how the wiring necessary for color vision can
develop; to do that, we need to know where PCS is first represented in the brain. By
"where is PCS first represented in the brain?" we mean, where is color represented in
a way that could be used directly by subsequent neural architecture to identify color?
We are looking in particular for some kind of reliable spatial/place coding of color.
Place or identity codes are not the only coding option, but they are the assumption of a

large body of neuroscience (e.g. see Arbib 2002), and therefore what we will focus on here and in our subsequent modeling work.

In this overview of the necessary background biology, we are focusing on the human and macaque early visual systems, which are generally similar (Preuss, 2004), and specifically similar for color vision (Jacobs, 2008). Much of our current understanding of vision comes from macaque studies. As Dacey (2000) reviews, macaques and humans have the same photoreceptor classes, with the same basic spectral sensitivities, and—so far—indistinguishable cell types and circuitry at least as far as V1. Macaques and humans also have similar visual abilities—behavior on color tasks is similar (e.g. De Valois et al., 1974; Stoughton et al., 2012). Here, we are discussing only daylight (photopic) color vision at light levels where color vision operates normally (i.e., we ignore photoreceptor bleaching, and the input of rods or other non-light-sensitive cells).

Very many experimental studies have investigated the representation of color in the early visual system, using a variety of techniques, from physiology through to psychophysics. Existing recent reviews include Solomon and Lennie (2007), Gegenfurtner and Kiper (2003), Conway (2009), and Shapley and Hawken (2011). Hence, we will not give comprehensive detail, but we will describe the three key stages of the early visual system from the point of view of the development of PCS.

In overview, color vision begins with light being absorbed by photoreceptors of the retina. The signals from these cones are then combined in a number of post-receptoral channels (RGC and LGN), allowing wavelength discrimination for the first time (a cone cannot distinguish radiance and wavelength). There does not appear to be any organization by color at the RGC or LGN stages. Following that, we have V1, where—unlike earlier stages—large-scale in vivo imaging is possible due to its location on the surface of the brain, and we start to see an organization by PCS. Figure 2.9 gives an overview of the structures involved, containing detail for each of the three stages that we will go on to describe in the following sections.

Before proceeding, it is important to note that the color vision pathway surely does not end at V1, nor is it likely to involve only a feed-forward pathway or a representation only in one particular location. In particular, V2 and V4 are typically implicated in color vision (see e.g. Conway, 2009 or any of the other reviews mentioned above). However, we choose to focus on the pathway up to and including V1 because V1 appears to be the earliest area of the brain in which any form of perceptually relevant organization for color appears, and we would like to determine (through modeling, in subsequent chapters) the minimum mechanism required to achieve such an organiza-

tion. We will return to this issue briefly in section 6.4.9.

2.5.1 Photoreceptors

As described earlier, human and macaque retina normally contains photoreceptors sensitive to three different bands of wavelength, termed the long (L), medium (M), and short (S) cones. Each receptor type is excited by a wide range of wavelengths, and the excitation also depends on radiance, as shown in figure 2.10. The cone sensitivities overlap; in particular, the L and M cones are highly correlated. In fact, L and M cones are morphologically similar, and their photopigments are genetically very similar, suggesting relatively recent evolutionary divergence (e.g. see review in Neitz and Neitz, 2011, or Jacobs, 2012). S cones are more distinct, physically and genetically.

Despite the typical picture of cone sensitivities (figure 2.10), it is important to note that different "color-normal" individuals have varied spectral sensitivities because of varying photopigments, varying density of photopigment in the cones, and differing pre-receptoral filtering e.g. from pigment in the lens and media of the eye (Stockman and Sharpe, 1999), such as the macular pigment. Macular pigment density also varies with eccentricity, and therefore the sensitivities for one individual vary with eccentricity. Macular pigment density can also vary between the two eyes. Additionally, sensitivity can change over time for one individual. For instance, the lens yellows with age, reducing short-wavelength sensitivity.

In both macaque and human, L and M cones appear to be arranged randomly or semi-randomly on the retina (Roorda and Williams, 1999; Roorda et al., 2001; Hofer et al., 2005a), and the ratio of L:M cones varies across the retina; M cones may even be absent after $8°$ (Kuchenbecker et al., 2008). The ratio of L:M cones also varies greatly between individuals (Roorda and Williams, 1999; Hofer et al., 2005a). S cones, however, are arranged more regularly, and comprise only 5–6 % of cones at $1°$ of eccentricity— the peak of their density, either side of which S cone density declines. S cones are nearly absent in the fovea, and are entirely absent in the foveola (leading to small-field tritanopia; Williams et al., 1981). Unlike for L and M cones, the distribution of S cone density is relatively consistent between individuals (Curcio et al., 1991; Hofer et al., 2005a). Figure 2.11 shows the layout of photoreceptors from several humans.

As mentioned in section 2.4.2, PCS is relatively constant between individuals, and for one individual across the visual field and over time. At this point—even only a short distance into the visual pathway—having reviewed several sources of physiological variety between individuals and within one individual, we should pause to review that
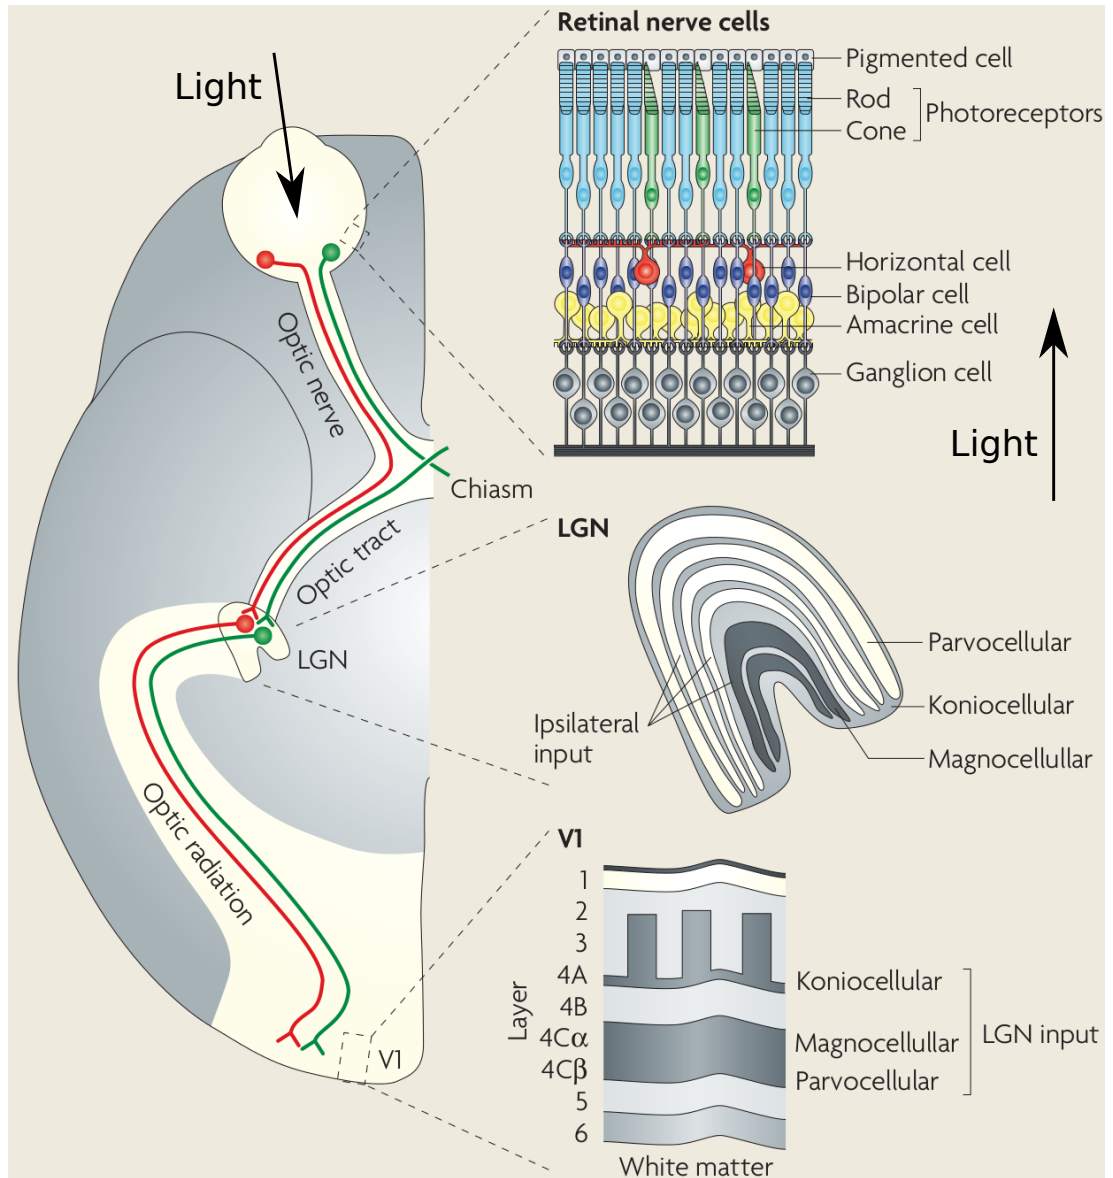
Figure 2.9: *Structures in the early visual system from the point of view of color vision. Left (approximately a horizontal plane through the brain): the retina is located at the back of eye, the LGN is located near the middle of the brain, and V1 is at the back, on the surface. The retina is a multi-layered sheet about $250\,\mu m$ thick. Light may be absorbed by one of three types of cone photoreceptor; ganglion cells ultimately compare the output of different cone types against each another, and the signal leaves the retina via the optic nerve. Different classes of ganglion cell project to different layers in the LGN. For color vision, we focus on P cells, which project to parvocellular LGN. Parvocellular LGN projects to V1 layer 4. Image reprinted with minor modification from Solomon and Lennie (2007).*
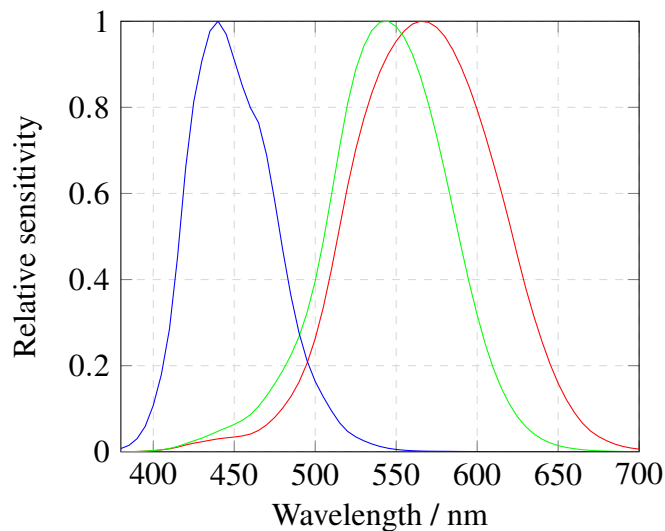
Figure 2.10: *LMS cone fundamentals. These Smith-Pokorny (Smith and Pokorny, 1975; Pokorny and Smith, 2013) LMS fundamentals show relative sensitivity at the cornea for human LMS cones (with each curve normalized to have peak 1.0). The functions are derived from human color-matching experiments; they are a linear transform of the CIEXYZ cmf. Measured sensitivities of the cone photopigments match these curves, when corrected for pre-receptor filtering (e.g. from the eye's media). Note that the cones may also loosely be referred to as "red", "green", and "blue", but this does not imply that e.g. "red" (L) cones are responsible for our perception of red. These spectra have provoked many questions. For instance, why the large overlaps (perhaps for spatial acuity; Lewis and Zhaoping, 2006)? Why have three types of cone, rather than more, or fewer (perhaps for finding fruit; Regan et al., 2001)? However, our focus is on ways in which this classic picture of cone sensitivities is potentially misleading. First, how should the curves be scaled? There are many more L and M cones than S cones. Furthermore, the cone responses adapt over both short and long timescales (see text): the cone responses are not given by static functions of the input. Also, the sensitivity curves vary between color-normal individuals because there are multiple variants (and densities) of each photopigment type (particularly L and M), and because pre-receptoral filtering varies (particularly at shorter wavelengths). Additionally, pre-receptoral filtering (and hence sensitivity) varies over time for an individual (e.g. an older lens may absorb far more short-wavelength light than a younger lens). Despite all this variability, we have a perceptual color space that is reasonably consistent between individuals, and for one individual over time.*
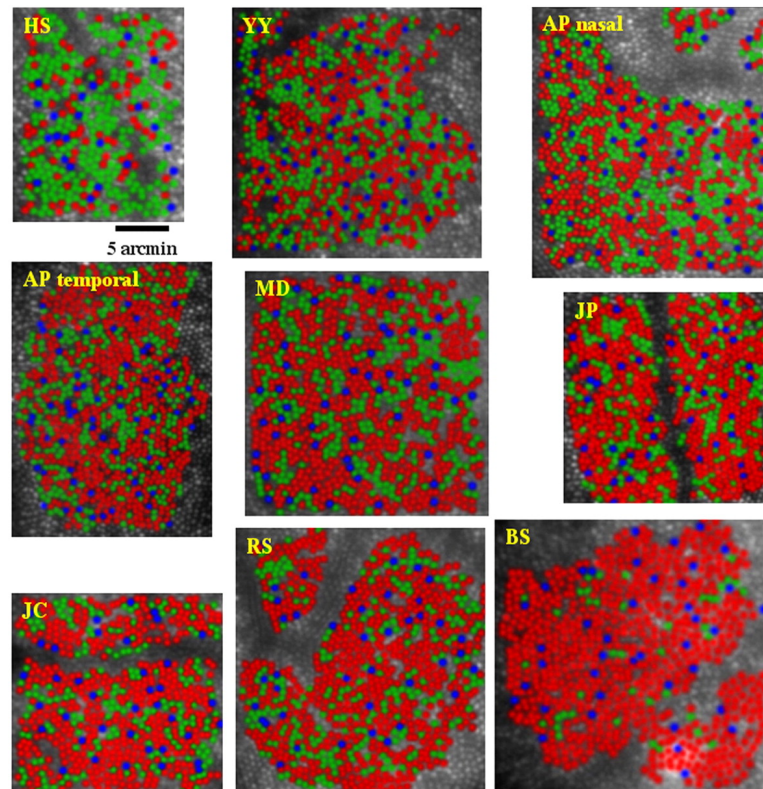
Figure 2.11: *Layout of cones. False color images (all to the same scale) of human L ("red"), M ("green"), and S ("blue") cones at $1°$ eccentricity. The ratio of L:M cones varies from 0.37:1 (HS) to 16.5:1 (BS), yet all subjects had normal color vision. L and M cones appear to be arranged somewhat randomly, while S cones are more regularly arranged. Data from macaque has shown similar properties (Roorda et al., 2001). It is still currently unclear whether there is a general bias towards humans having more L than M cones, with the ratio being more equal in macaque. It is also not clear if the arrangement of L and M cones is truly random in both species. Images reprinted from Hofer et al. (2005a).*

claim. First, for different individuals:

- Different individuals have a wide variety of spectral sensitivities (for several reasons, described earlier), but differences between their sensitivities do not predict what stimuli observers call white or a particular hue (Webster et al., 2000; Neitz et al., 2002).

- Different individuals have highly variable ratios of L and M cones, apparently randomly (or semi-randomly) arranged on the retina. Despite this, observers agree on the wavelength of "unique yellow" (appearing neither reddish nor greenish) to within a few nanometers (Brainard et al., 2000; Neitz et al., 2002; Miyahara et al., 1998).

- Diabetics with low S-cone sensitivity make judgments of unique hues within the normal ranges (Schefrin et al., 1991).

- There are examples of individuals with quite abnormal physiology, yet who are able to make reasonably normal color judgments. For instance, an individual with little cone function, abnormal macular pigment, and no L cone pigment was able to make unique blue and green hue settings within the range of color-normal observers, and unique yellow and red only slightly outside the normal range (Crognale et al., 2001).

Second, for an individual:

- Spectral sensitivity varies across the center of gaze (primarily because of variations in macular pigment density), but the color appearance of large uniform fields remains constant (O'Neil and Webster, 2014), and colors (including white) shown at various eccentricities are judged to be far more similar than the different sensitivities would predict (Webster et al., 2010; Beer et al., 2005). This compensation is not perfect, however, and may vary between hues (see e.g. Parry et al. 2006; Bompas et al. 2013; O'Neil and Webster 2014).

- As stated above, the appearance of color remains constant across a wide range of visual field, despite changes in effective cone sensitivity. Similarly, the arrangement and ratio of cones varies across the visual field, yet this also does not seem to have much effect. For instance, the center of the fovea (the foveola) contains only L and M cones (i.e. no S cones), and it is possible to demonstrate inability to detect short-wavelength light over this region (Williams et al., 1981)—yet we do not perceive this except under specific, unusual, situations (Magnussen et al., 2001).

- After cataract surgery, suddenly far more short-wavelength light reaches the retina than over many preceding years, and color judgments are strongly changed. For instance, the stimulus judged as achromatic after cataract removal is shifted towards yellow relative to its location before surgery. However, over a period of months, the point shifts back towards its original location (Delahunt et al., 2004).

- Aging causes large changes in sensitivity (e.g. an increasingly yellow lens), but this has a lower than predicted impact on color vision (Werner et al., 2004). The stimulus judged as achromatic does not significantly differ between individuals of a wide range of ages (Werner and Schefrin, 1993). Additionally, there is also no significant difference in some unique hues, although as with change in sensitivities with eccentricity, this may vary by hue (Schefrin and Werner, 1990; Wuerger, 2013).

How can a relatively stable and consistent PCS be built on such unstable/inconsistent equipment? Alongside the important signal (about objects in the world), each level of the visual system also encounters irrelevant variabilities (ones not important to survival) from the stage before, starting with relatively unimportant variabilities in the incoming light—as discussed in section 2.4.3. Chromatic adaptation must therefore be dealing with physiological variability to a large degree, too. The actual mechanisms and locations underlying chromatic adaptation are not certain, but it seems clear that some adaptation takes place at each of the three key stages of our early visual system pathway, starting with the photoreceptors.

2.5.1.1   Gain control and adaptation

Earlier we described the sensitivities of the cone classes. Cone excitation is linear; that is, if radiance is doubled, the (incoming) excitation of a cone is also doubled. However, an important property of the cones is that their (outgoing) *responses* (activities) are non-linear. While the exact mechanism is not yet clear, photoreceptor response rates do appear to adapt on short timescales to accommodate the incoming range; that is, the output of a cone depends on its gain as well as its spectral sensitivity (see figure 2.12). Furthermore, the effective output from the cones may be modified by horizontal cell connections, allowing a further spatial element to the normalization of photoreceptor responses. In any case, short-term gain changes could provide the mechanism for a von Kries-like chromatic adaptation.

Secondly, cones adapt over longer timescales. Again, there is no clear mechanism, but
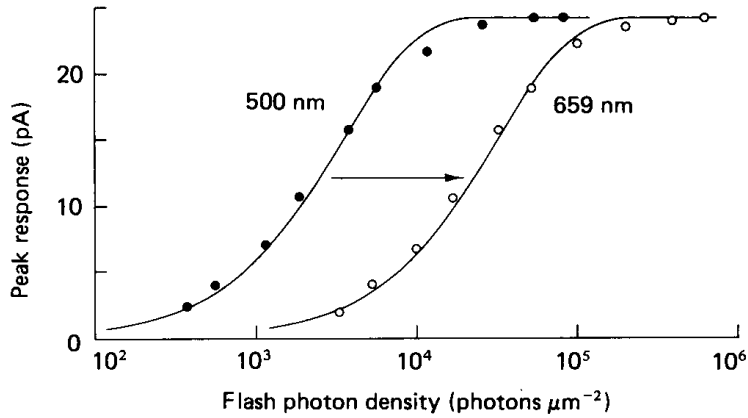
(a)
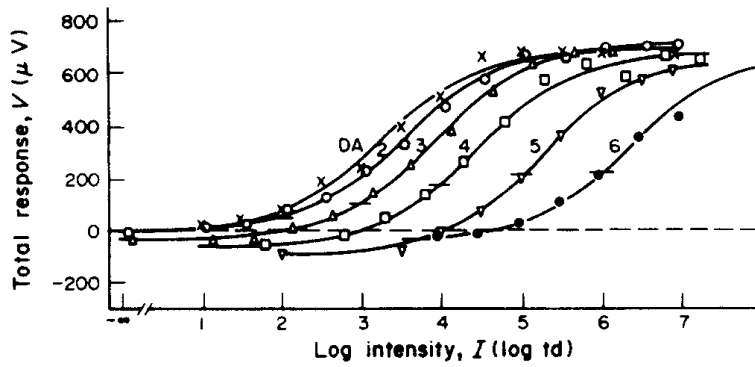
(b)

Figure 2.12: *Cone photoreceptor response. (a)As shown earlier in figure 2.10, the wavelength of incident light affects a cone's response (here photocurrent of a macaque cone is being measured). However, (b) shows that short-term adaptation also has a significant effect on cone responses (here from rhesus monkey). The nonlinear cone response can be modeled by an equation of the form $\frac{V}{V_m} = \frac{I^n}{I^n + \sigma^n}$ (Naka and Rushton, 1966), where σ describes the adaptation, including the cone's dark adapted state, its level of "neural adaptation" (mechanism unclear), and its level of photopigment bleaching. Image in (a) reprinted from Baylor et al. (1987); image in (b) reprinted from Valeton and van Norren (1983).*

possibilities include changes in outer segment length (found in fish) and horizontal cell connectivity changes, effectively modifying the output of cones (Wagner and Kröger, 2005). It is also not clear what the long-term adaptation state of cones would be, but one suggestion is that cones are adapted to long-term "white", i.e. meaning all cones would respond equally to the average daylight spectrum (Walraven and Werner, 1991).

Note that the short- and long- term adaptation of the cones means the diagram of cone sensitivities in figure 2.10 is misleading, because it presents a static picture at one state of adaptation, when the reality is highly dynamic.

2.5.1.2   Color representation

Finally, we consider color representation in the photoreceptors. The first thing to note here is that the output of a single photoreceptor cannot be used to distinguish wavelength from intensity. That is, any given response level for a photoreceptor could represent a certain intensity and monochromatic wavelength, or e.g. a less effective wavelength at a higher intensity (figure 2.10). Only by comparing photoreceptors of different types can color vision begin. Other evidence against photoreceptors providing an obvious representation of perceived color includes that changing photoreceptor absorption ratios does not necessarily change perceived hue: broadening the spectrum of light that is already giving a certain hue percept does not change the percept (Mizokami et al., 2006). Additionally, making it clear that a particular cone type is not responsible for a particular hue percept, Hofer et al. (2005b) demonstrated that stimulation of one cone class can elicit multiple different hue percepts, including that seeing "blue" colors does not require S cones.

In this section about the photoreceptors, we have emphasized variabilities in physiology, adaptation of responses, and PCS not being represented. As noted above, the isolated signal from a single cone cannot even be used to distinguish between wavelength and intensity. However, a comparison of the signals from different cone classes at nearby locations on the retina—cone opponency—does allow such a distinction. Such a comparison is performed at the next stage of our pathway, the RGC.

2.5.2   RGC and LGN

Figure 2.9 shows that after the photoreceptors, there are three more classes of cell (horizontal, bipolar, amacrine) before the ganglion cells, which project via the optic nerve to the LGN (near the center of the brain). We will not explicitly describe the
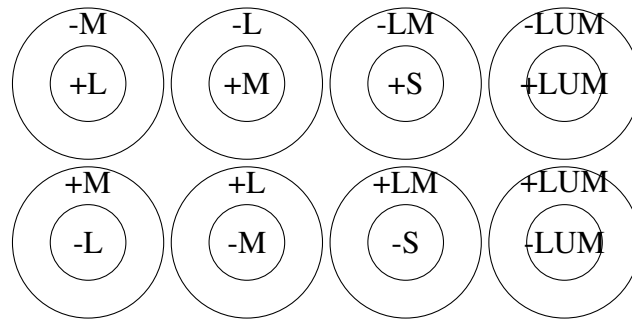
Figure 2.13: *Spatial and cone opponency in the RGC and LGN . Top left: By comparing the output of an L photoreceptor ("L center") with that of M photoreceptors ("M surround"), wavelength and intensity changes at that location can be distinguished. This cell is excited by "long-wavelength" (L) light in the center, and inhibited by "medium-wavelength" (M) light in the surround, so it is referred to as L+/M- ("L ON center, M OFF surround"). Seven other generally acknowledged chromatic ganglion cell classes are additionally shown. The S/LM pathway may have a different spatial structure, perhaps being coextensive rather than having a center/surround arrangement, but recent studies have found evidence of both (Field et al., 2007; Crook et al., 2009). In fact, despite the classical picture, exactly what classes of retinal ganglion cell exist (cone-opponent and otherwise) is still unclear (Neitz and Neitz, 2008), and it seems possible there is a wider variety than the classes shown above. Such variety may result from "random wiring" of a random cone mosaic (Lennie et al., 1991). For the parafoveal pathway (on which subsequently reviewed modeling focuses), ganglion cells generally have a center comprising one cone, but a surround comprising approximately 5 cones.*

horizontal, bipolar, and amacrine cells, instead assuming their combined effects are summed up in the behavior of the ganglion cells (since only the ganglion cells project to the brain). There are two major classes of ganglion cell that project to the LGN, the P cells and the M cells. We will provide more detail in a subsequent section, but the P cells are of most interest for color vision (e.g. Lee et al. 1990), and are the most numerous ganglion cell type. They project to parvocellular cells in the LGN, which have similar properties.

The P ganglion cells represent the next stage in the creation of color by the brain: they compare signals from different cone classes at nearby locations on the retina. Macaque RGC seems to contain three main classes of cone-opponent P ganglion cells, generally spatially and/or cone opponent: luminance (i.e. L+M), L vs. M, and S vs. (L+M) (Dacey, 2000; figure 2.13). However, the number of classes, and exactly what each class is, is not clear (Neitz and Neitz, 2008). It seems likely that the cone connections, particularly to L and M cones forming the surrounds, are not as specific as in the textbook presentation (see e.g. Field et al. 2010), and therefore perhaps there are varying degrees of cone opponency. We will revisit this possible "random wiring" in chapter 6, but for now we focus on the classic picture shown in figure 2.13.

There are various possible explanations for why we have cone opponency. Cone-opponent cells directly give information about spectral content, unlike the cones (e.g. increased L cone response could merely be from an increase in irradiance, whereas increased L-M response must be from an increase in L cone wavelengths vs. M cone wavelengths). Also, L and M retinal ganglion cells may be spectrally sharpened vs. cone responses (Foster, 2010), an important component of color constancy (Finlayson et al., 1994; Vazquez-Corral and Bertalmío, 2014). Cone opponency also reduces information redundancy (Buchsbaum and Gottschalk, 1983; Atick et al., 1992; Dan et al., 1996): L and M cone responses are highly correlated, whereas L-M and M-L responses are less so (Zaidi, 1997), and at least some of the reduction comes from opponent processing (Pitkow and Meister, 2012). There are various reasons this could be important, including the limited capacity of the optic nerve and limited dynamic range of spiking neurons. Some studies have found that in natural scenes, chromatic information varies independently along these two cone-opponent axes (e.g. Ruderman et al. 1998) for certain types of scene (e.g. foliage-dominated scenes). Opponent processing in general is reviewed in Westheimer (2007).

2.5.2.1   Gain control and adaptation

As with the photoreceptors, determining the extent of gain control and adaptation in retinal ganglion P (midget) cells has proven difficult. Additionally, since retinal ganglion cells inherit the properties of photoreceptor responses, it can be difficult to determine whether observed behavior originates at the photoreceptor or ganglion cell level. And again, as for the photoreceptors, ganglion cell responses may adapt because of intrinsic changes to the cells themselves, and/or because of the effects of connections (e.g. involving amacrine and horizontal cells). While adaptation in P cells appears to be less than in M cells, P cells do show short-term adaptation (Kaplan and Benardete, 2001; Lee et al., 2008; Dunn et al., 2007; Yeh et al., 1996), and if photoreceptors compute contrast, then adaptation in the ganglion cells can give contrast adaptation (since they receive input from the photoreceptors). However, it is not at all clear whether this might translate into any form of chromatic contrast gain control. It is also unclear whether P ganglion cells adapt over long timescales, but they do at least inherit the adaptation of the photoreceptors, and have been shown to adapt over timescales of 10–100 s (McLelland et al., 2009).

2.5.2.2   Color representation

Although the retinal ganglion cells and LGN contain a number of cone-opponent channels, the reported receptive fields do not correspond to the large number of the hues we experience (Neitz and Neitz, 2008). As reviewed above, RGC and LGN responses are generally considered to fall into just the categories L vs. M and S vs. (L+M) (as well as a luminance response, L+M), so shown in figure 2.14. Ganglion cells therefore perform wavelength discrimination, but it is not clear that they provide coding of perceptual color. There is also no evidence of spatial organization for color; neither retinal ganglion cells nor LGN cells appear to be organized by perceptual color.

While not yet certain, recent work (to be described in the following sections) indicates that V1 is the earliest area of the brain from which signals can be detected corresponding to perceptual hues, and hence it is the focus of this thesis. Unlike the retinal ganglion cells and LGN, V1 is accessible to large-scale imaging in living animals, and results from imaging studies indicate that V1 shows some organization of color responses by PCS, as will be described in the next section.
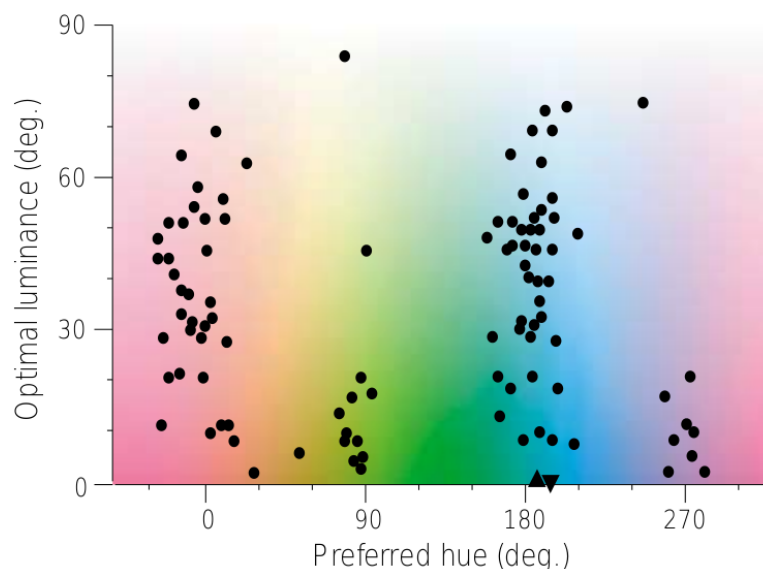
Figure 2.14: *Optimal colors for driving LGN neurons. Most LGN cells, like retinal ganglion cells, cluster into a few categories. Note: these colors are defined on physiological (DKL) axes. Image reprinted from Gegenfurtner (2003); data from Derrington et al. (1984).*

2.5.3   Primary visual cortex

Input from parvocellular LGN arrives in layers 4Cβ and 4A of macaque V1. V1 is well known to have interesting spatial organizations of various features that would support their coding. For instance, neurons in cat V1 were discovered to have tuning for OR (Hubel and Wiesel, 1959), and subsequently with optical imaging (OI) and related techniques it has been possible to see that there is a retinotopic organization into a spatially contiguous map of OR preference. Such a map has been found in macaque V1 (Blasdel, 1992b), along with maps of other features such as ocular dominance (Blasdel, 1992a). These maps are just two ways of analyzing the same sheet of neurons, and many such maps allow the whole range of multiple features to be represented together. Typically, the maps appear to be optimized for coverage (Swindale et al., 2000; Swindale, 2004), and are typically also continuous (Blasdel, 1992b; Yu et al., 2005).

However, equivalent results for CR have been more controversial. The earliest studies using electrophysiology found OR and CR were largely segregated in V1: CR responses appeared to be confined to small, spatially isolated regions—in particular, corresponding to the mitochondria-rich cytochrome oxidase (CO) blobs (Hubel and Wiesel, 1968; Livingstone and Hubel, 1984)—leading to the view that OR and CR are processed separately in V1 (Livingstone and Hubel, 1988; Gegenfurtner, 2001). More

recent electrophysiological studies have continued to find such separation (e.g. Con-way and Livingstone, 2006), but at the same time other electrophysiological studies have found OR and CR preferences in macaque V1 to be overlapping, with a large number of cell types with varying achromatic and/or chromatic tunings (Lennie et al., 1990; Johnson et al., 2001; Friedman et al., 2003; Johnson et al., 2008; Wachtler et al., 2003). While electrophysiology can determine the properties of single cells (and indeed gave the first indication that neurons are preferentially sensitive to a small set of features, and that preferences varied across space), it is difficult to sample a large enough number of cells and to reconstruct their positions accurately across a large enough area of the cortex (Olshausen and Field, 2005). On the other hand, larger-scale techniques such as functional magnetic resonance imaging (fMRI) suffer from the opposite problem: the resolution is too coarse, and this could be a problem for hue in particular because of the apparently smaller scale of hue organization in macaque V1 (with perhaps the most important processing being confined to patches, as we will see).

Following on from electrophysiological studies, imaging studies using color and luminance stimuli in V1 have been able to determine the responses of a large enough population of cells to allow conclusions to be drawn about map-level responses, while also having a high enough resolution (which, for in vivo two-photon calcium microscopy (2P), is at the level of single cells). Imaging studies so far have shown that while orientation-selective cells form contiguous maps in V1, CR cells appear to cluster into patches. Furthermore, the organization of cells within patches appears to be related to perceptual color. As was introduced in figure 1.9 in the previous chapter, an optical imaging study in macaque (Xiao et al., 2007) shows the peak responses to different hues progressing following the same pattern in V1 as in perceptual color space. Subsequent preliminary 2P studies support this finding. We will review this work in detail in the following sections.

2.6   Representation of CR and OR in V1

Having introduced V1 as an important target for understanding the coding of OR and CR, and the debate surrounding the extent of CR/OR separation and the representation of PCS, in this section we will review results from experimental studies about the representation of OR and CR in macaque V1 in more detail.

We are focusing on experimental results about joint OR and CR organization, and on the organization of hue, in macaque V1. Therefore, here we do not cover results from

	Chatterjee (2010)	Lu and Roe (2008)	Xiao et al. (2007)	Landisman and Ts'o (2002)	Blasdel (1992b)
Species	(macaque)	Macaca fascicularis	Macaca fascicularis	Macaca fascicularis & mulatta	Macaca nemestrina
Eccentricity	parafoveal	parafoveal $(1.5° - 5°)$	parafoveal $(2° - 4°)$	foveal, perifoveal, & $5° - 10°$	?
Imaging	2P	iOI	iOI	iOI	iOI
Features imaged	OR, CR	OR, CR	CR	OR, CR, OD	OR

Table 2.1: *Key features of macaque V1 OR/CR studies. We are considering only macaque studies, although the precise species varies between studies. Studies also do not image exactly the same region of V1, and they do not all present the same stimuli or image the same features. Other undocumented but potentially important factors may differ, too. For instance, what was the rearing environment of the animals used in each study? Were all the animals checked for normal luminance and color vision? Are the results an unbiased survey, or are only specifically selected results presented (potentially reducing the variety of results from each study)?*

other other species, other areas of the macaque visual system, or studies that do not investigate OR or CR representation. Also, we focus on imaging studies because, as mentioned earlier, they are large enough scale to be behaviorally relevant while also having enough resolution. The studies we will be using to build a picture of V1 OR and CR organization are summarized in table 2.1. All investigate macaque V1, but they differ on the exact species, V1 imaging location, stimuli used, and other attributes. We will discuss similarities and differences as we go along.

2.6.1 Spatial organization of OR: OR maps

In macaque V1, as with other species such as cat and ferret, it is widely accepted that there are spatially contiguous maps of OR preference. Figure 2.15 shows a typical OR preference map, indicating many of the features usually found (such as iso-orientation zones and pinwheels). These maps have a ring-shaped Fourier transform indicating the frequency with which orientations repeat (on average, in all directions). The fine structure of OR maps is supported by 2P results e.g. in cat, showing the detailed cell-level makeup of the map (Ohki et al., 2006).

Given the prevalence of OR maps (and indeed other topographical maps) in mammals,

we might ask what (if anything) maps are useful for in general (Horton and Adams, 2005). For instance, perhaps maps minimize wiring length by allowing connected neurons to be as close together as possible (Chklovskii and Koulakov, 2004)? However, the present review, and our modeling work in later chapters, does not require the patterns to serve any functional purpose: we are asking what organization is found (if any), and how such organization might arise. We can see that OR maps measured in macaque have certain features (e.g. their spatial organization, smoothness, and so on), and a model of macaque V1 therefore ought to match these—but it is the reliable association between location in the map and OR that is key. This type of coding—place or identity coding—is not the only coding option, but it is an assumption of a large body of neuroscience (Arbib, 2002). Indeed, some mammals do not have OR maps with the same kind of spatial organization as in macaque, but the reliable association between cortical position and orientation tuning is still important. Mice, for instance, have an OR representation that does not appear to be spatially organized, but cells are still reliably tuned to particular orientations (Ohki and Reid, 2007). For a review of the purpose of topological maps, arguing that we cannot conclude the spatial patterns have a computational role vs. being an epiphenomenon of cortical evolution, see Wilson and Bednar (2015).

2.6.2   Spatial organization of CR 1: CR patches

In macaque monkey, several OI studies indicate that color-selective cells in V1 occur in isolated patches, rather than forming a spatially contiguous map as is the case for orientation-selective cells. In an iOI study of macaque parafoveal V1, Landisman and Ts'o (2002) subtracted the responses to high-contrast luminance (black/white; BW) gratings from the responses to red/green (RG) isoluminant gratings (all stimuli equiluminant on average), finding a patchy distribution of CR cells and a contiguous OR map (figure 2.16). The OR map agrees with previous studies of OR representation in V1 (e.g. Blasdel 1992b).

Again using iOI of macaque V1, Lu and Roe (2008) found patchy responses to color stimuli (and a contiguous OR map), in agreement with Landisman and Ts'o (2002). Again, the response to BW gratings was subtracted from the response to isoluminant RG gratings, and the resulting differential images show patches where the response to RG gratings is higher than to BW gratings (figure 2.17A). In addition to showing the organization of CR cells into patches, Lu and Roe (2008) also show that the OR selectivity of the CR patches is lower than that of randomly chosen locations (figure 2.17D).
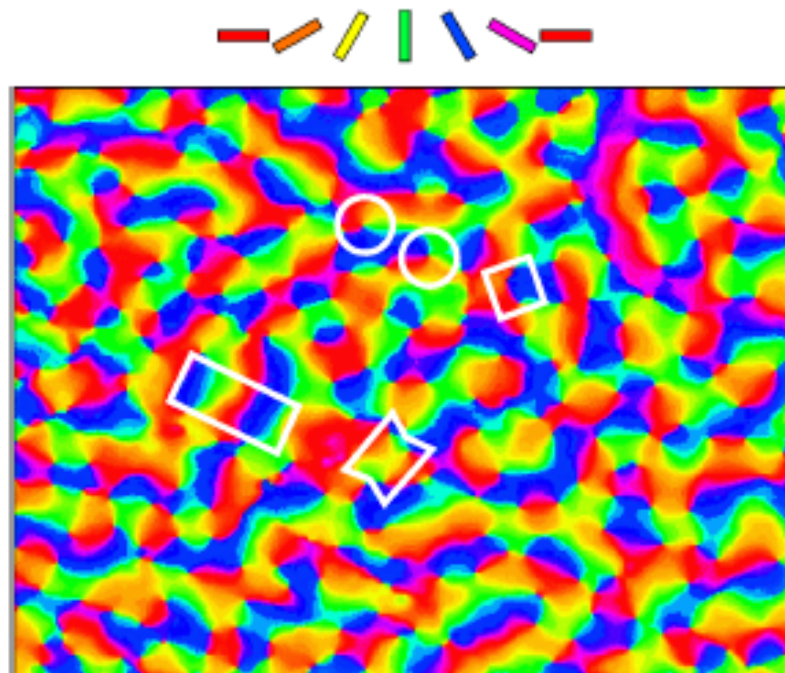
Figure 2.15: *Macaque orientation preference map (Blasdel, 1992b). By presenting drifting gratings of multiple orientations, neurons' preferred orientation can be measured using OI. Each pixel in the map is colored according to orientation preference, and represents the average of about 1000 neurons. Nearby neurons generally have similar preferences, resulting in iso-orientation zones. Other features typically found are pinwheels (points around which all orientations are represented continuously; circles), fractures (sudden changes in preference; square), linear zones (smoothly changing preferences; rectangle), and saddle points (partial bisection of one iso-orientation zone with another).* $7.5\,\mathrm{mm} \times 5.5\,\mathrm{mm}$ *region of parafoveal V1. Image reprinted from Blasdel (1992b).*
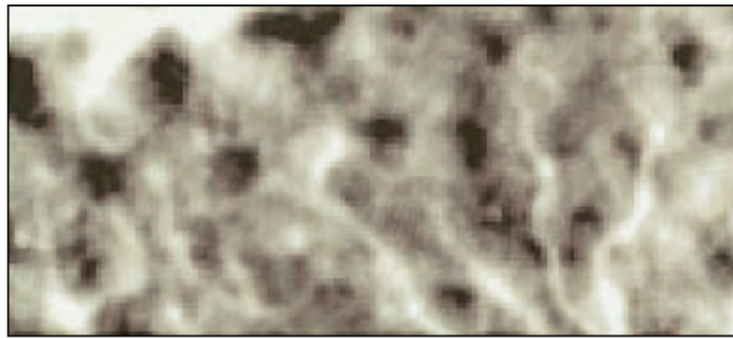
(a) Color selectivity map
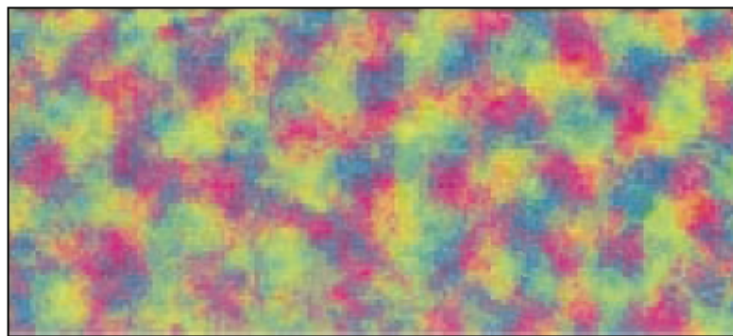
(b) Orientation preference map

Figure 2.16: *CR patches and OR maps (Landisman and Ts'o, 2002). While OR appears to be represented in spatially contiguous preference maps, CR appears to be represented in patches.* $5.3\,\text{mm} \times 2.3\,\text{mm}$ *region of macaque parafoveal V1 from iOI. Images reprinted from Landisman and Ts'o (2002).*
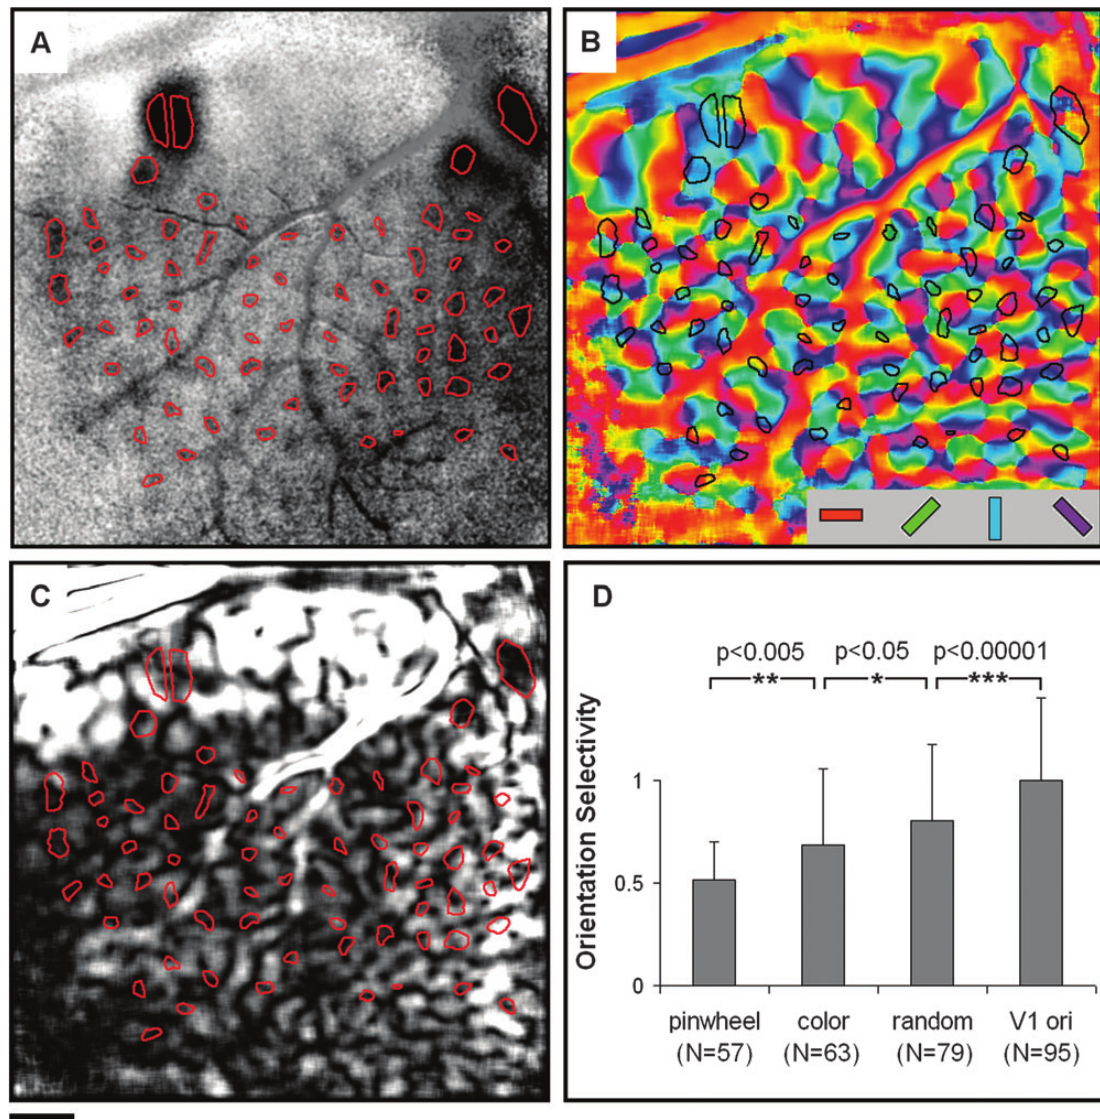
Figure 2.17: *CR patches and OR maps (Lu and Roe, 2008). Color responses occur in patches, and those CR patches occur in regions of low OR selectivity. A: Response to LUM (achromatic) gratings subtracted from response to red/green isoluminant gratings, showing responses that are higher to CR than to LUM occurring in patches. Red circles are drawn at a threshold of the most active 10% of pixels. B: Orientation preference map with CR circles from A overlaid in black. C: Orientation selectivity map, again with CR circles overlaid. D: Orientation selectivity in different regions, indicating that color patches have lower selectivity than other areas (*). Scale bar for A–C 1 mm. Images reprinted from Lu and Roe (2008).*
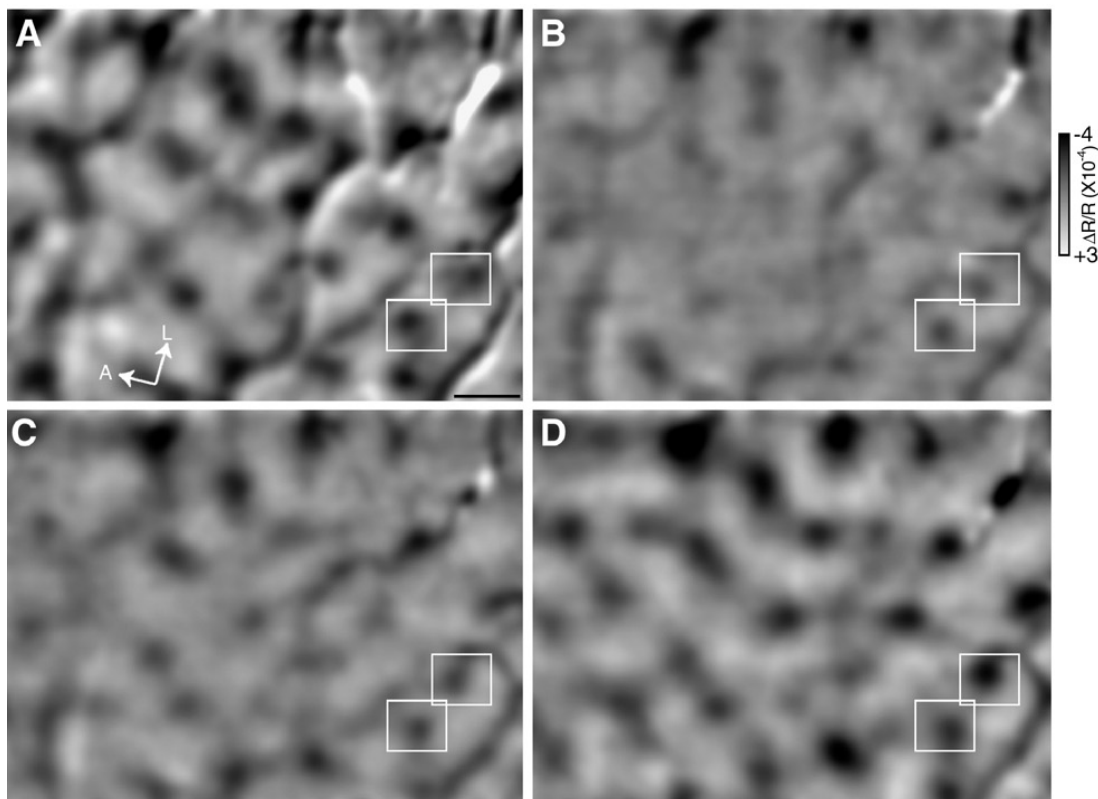
Figure 2.18: *CR patches (Xiao et al., 2007). V1 responses to uniform stim-
uli of different colors (A–D: red, yellow, green, blue). Patches correspond
to those identified by RG-BW gratings (not shown). Scale bar* $500\,\mu m$*. Re-
gions in white rectangles will be expanded in the following section. Images
reprinted from Xiao et al. (2007).*

Xiao et al. (2007) identified CR patches in the same way as earlier studies: for 0
and $90°$ degree orientations, responses to achromatic (BW) gratings were subtracted
from responses to chromatic gratings (isoluminant RG gratings). However, this study
differs from the previous two by (a) not also measuring an OR preference map and (b)
additionally presenting spatially uniform, photometrically isoluminant color stimuli.
The responses to both RG-BW gratings and spatially uniform stimuli were patchy, and
although there is no detailed analysis, the patches from both methods appear to align.
For each single hue presented, significantly activated regions (vs. gray) were identified.
The patchy responses to each hue combined to give patches approximately $200\,\mu m$ at
half height, typically separated by more than $400\,\mu m$ (figure 2.18).

Preliminary results from in vivo 2P of macaque parafoveal V1—described in Chatter-
jee et al. (2008), Ohki et al. (2008), Chatterjee (2010), and Conway et al. (2010)—are
consistent with the results from optical imaging. Spatially uniform color flashes cause
patchy responses, whereas responses to colored gratings (e.g. red/gray—not isolumi-

nant) do not cluster in patches (figure 2.19). The spatial scale of patches matches the OI results. Additionally, the neurons in CR patches are less selective for orientation than neurons outside the patches, and neurons in CR patches were rarely achromatic. However, none of the 2P findings have been published fully yet.

In summary, then, experimental data is broadly consistent and shows CR organized into patches alongside a spatially contiguous OR preference map. There are differences in CR patch sizes and spacing, but the studies have some methodological differences which could account for this (table 2.1). The results are scaled together and summarized in figure 2.20.
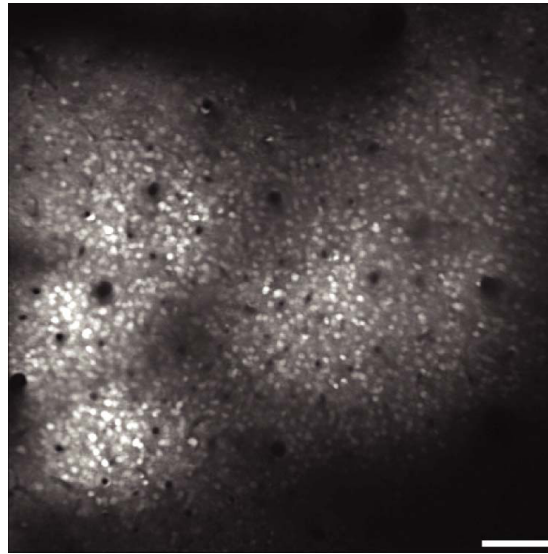
(a)
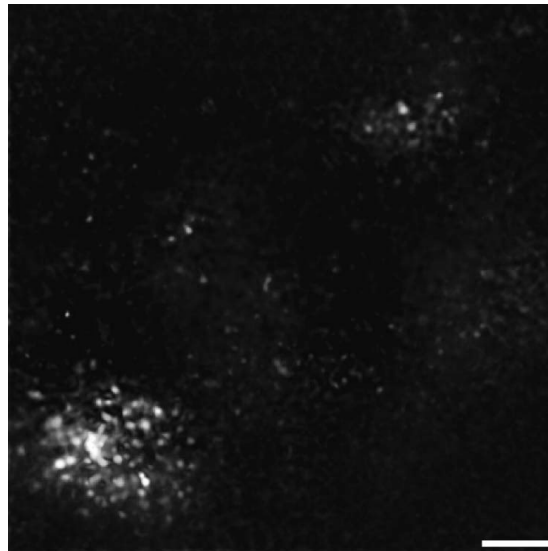
(b)

Figure 2.19: *CR patches (Chatterjee, 2010). Preliminary 2P study. (a) Positions of cells. (b) Combined responses to spatially uniform isoluminant color flashes of various L, M, and S cone contrasts (preceded and followed by gray), showing color responses in patches. Regions between the patches, which here do not respond to spatially uniform color flashes, did respond to colored gratings (e.g. red/gray). The organization of color responses into patches is consistent with earlier results from iOI. Scale bars are $100 \, \mu$m. Figures from Chatterjee (2010); (b) converted to grayscale.*

Figure 2.20 *(following page)*: *Comparison of experimental data: CR patches in OR maps. The first three columns are studies specifically investigating the representation of CR, while the final column shows a typical OR map. Each 1.5-inch square is $3\,\mathrm{mm} \times 3\,\mathrm{mm}$ of macaque V1. NA indicates no data available from that study for a particular condition. First row: differential image produced by subtracting response to black/white (BW) gratings from response to red/green (RG) isoluminant gratings. Scale is white (lowest response to CR) to black (highest response to CR). Red outlines are drawn at a threshold (we drew outlines on Xiao et al., 2007 and Landisman and Ts'o, 2002) and are used in subsequent rows to denote the CR patches. Second row: OR selectivity map. Note that OR selectivity is not patchy, and that CR regions tend to occur in areas of low OR selectivity. Third row: OR preference map. Fourth row: combined OR preference and selectivity map. Results are generally consistent: CR is organized into patches, although details are not consistent (e.g. different inter–CR-blob spacings relative to the OR map, different CR patch shapes).*
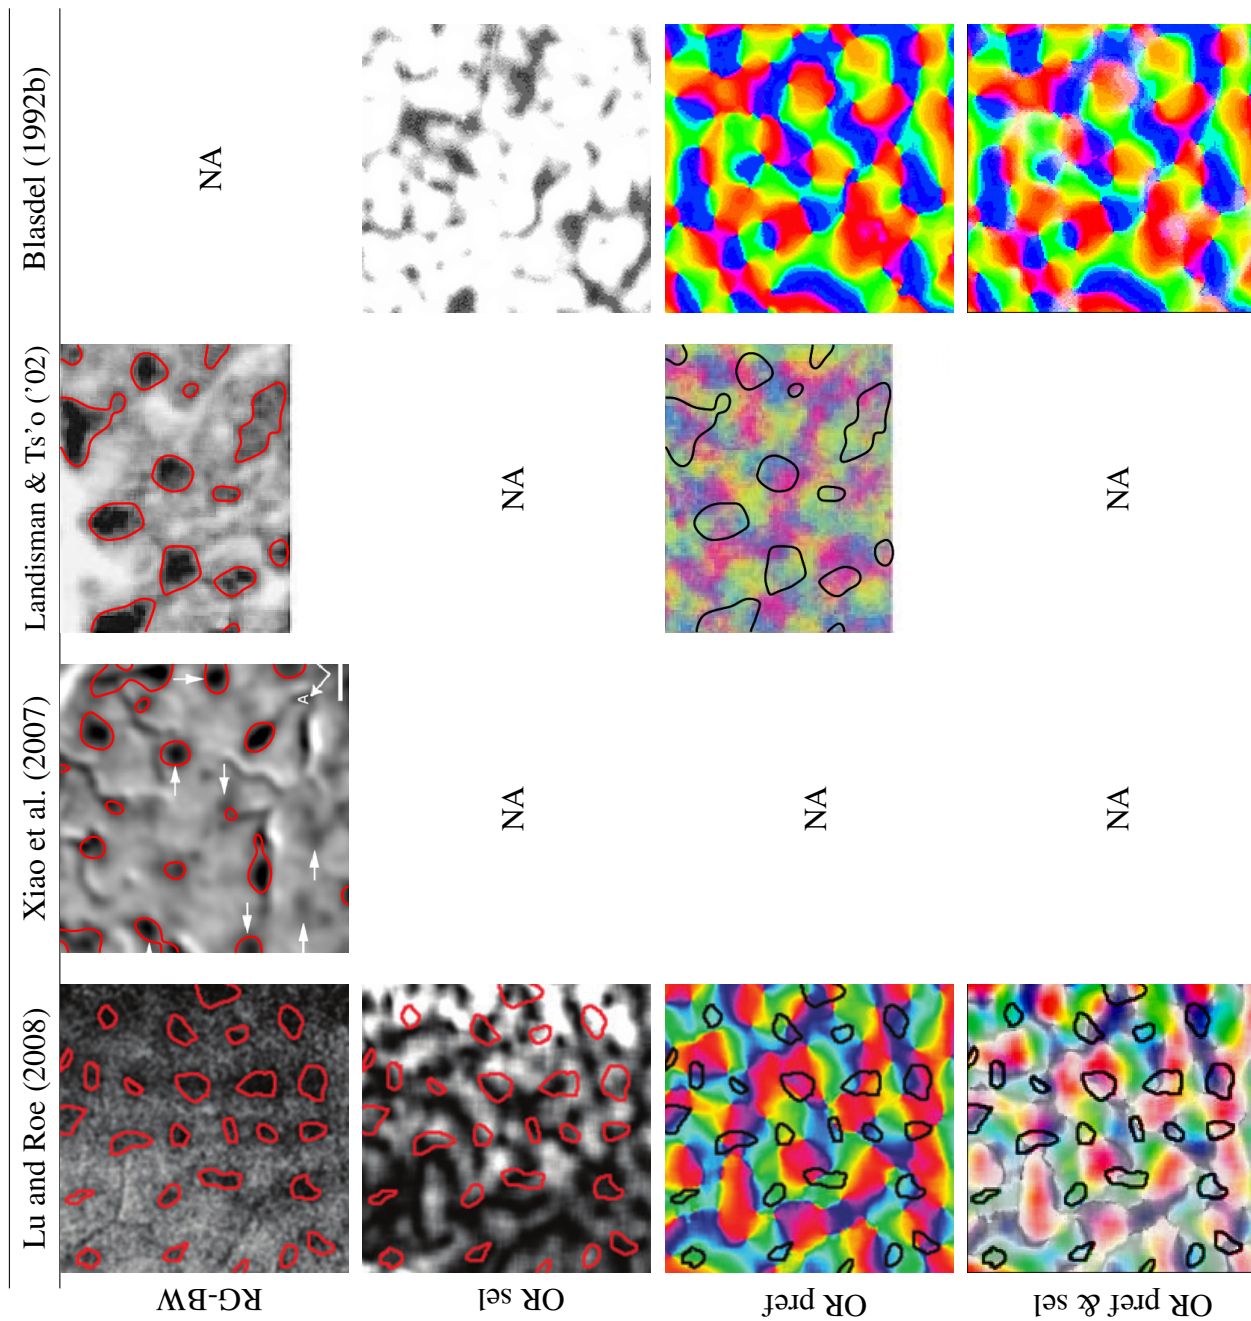
2.6.3 Spatial organization of CR 2: Hue maps

We have seen that in contrast to OR, CR appears to be represented primarily in spatially isolated patches. These patches tend to occur in regions of low OR selectivity. While the experimental studies broadly agree about joint OR and CR organization, the organization of perceptual attributes of CR such as hue and saturation is less clear. Here we focus on the organization of hue. The number of studies addressing the organization of hue in V1 by OI or 2P techniques is small. However, we believe it is worth focusing on these techniques because the results from techniques such as single-cell electrophysiology can be very confusing (it is difficult to sample enough cells and reliably reconstruct their positions), while the results from fMRI are at too low a resolution.

As for OR, we are looking for a representation of hue that allows coding by position. For instance, a place code for hue—as we will see suggested by Xiao et al. (2007) below—would allow neurons above V1 to reliably distinguish hue as represented in V1. Such a place code is not the only coding option, and indeed the spatial organization is not critical for coding, but if we find that cells are organized in this particular way in macaque V1, this is relevant for modeling. We will see below that experimental data tentatively indicates that each CR patch responds to a wide range of (or all) hues, and within a patch there is spatial organization according to perceptual hue.

The first indication that an organization for hue is present in CR patches comes from Xiao et al. (2007). We introduced this study in the previous section, mentioning that in addition to imaging CR patches with RG-BW gratings, spatially uniform stimuli of different hues were also used. Presenting spatially uniform stimuli of different hues in conjunction with imaging the patches at higher resolution would allow a type of "hue map" to be measured; that is, within the patches, the peak response to each hue varies in spatial position, and additionally, distances between peaks correspond to distances between hues in perceptual space. Since we are particularly interested in the spatial organization of hue, we will describe the method used in more detail.

Each color stimulus was presented after gray, and hues were selected from CIELUV so that hue differences would be perceptually uniform. As described in the previous section, for each single hue presented, significantly activated regions vs. gray were identified, giving patches approximately $200\,\mu m$ at half height, typically separated by more than $400\,\mu m$ (figure 2.21), and corresponding to the patches identified previously using RG-BW gratings. Within each region, the peak response to each hue was noted, along with a contour at 75 % of the peak. Responses to different hues overlap, but the peaks are spatially separated, and the distances between peaks correlate with distances

in PCS (figure 2.21(b)).

The analysis is only of the peak responses (and response contours) to different hues; there is no actual preference map like those shown for OR. Still, the correlation of the spatial arrangement of peak responses and perceptual color differences is exciting. The authors do note that responses to "end-spectral hues" (red and blue) were stronger than those to mid-spectral ones (green and yellow), so e.g. the spatial location with the peak response to yellow actually responded more strongly to red than to yellow. We can therefore speculate that a preference map constructed from this data might look end-spectral colored only. However, while the stimuli were photometrically isoluminant, they were not necessarily all identical in every way except for hue (e.g. the saturation of stimuli may have varied), which could obscure an actual organization by hue preference. The "color properties" of the stimuli to use for map measurement is a difficult question (for instance, Landisman and Ts'o 2002 were unable to obtain differential OI responses at all to their BY-BW gratings), and we will return to it in subsequent chapters.

The preliminary results from in vivo 2P of macaque parafoveal V1 introduced in the previous section are consistent with the results from optical imaging. The spatially uniform color flashes that highlighted color patches also showed responses whose positions depended on color (figure 2.22). This is important, because a criticism of OI is that its resolution may not be high enough compared to electrophysiology, and therefore results about hue responses within CR patches may not be trustworthy. The preliminary 2P studies have used stimuli with varied L/M/S cone contrasts, rather than stimuli defined by perceptual hue, so the CR patches were found to be organized according to L/M/S contrast. While not the same as being organized by perceptual hue, this is not inconsistent with being organized by perceptual hue.

In summary, Xiao et al. (2007) report that each patch responds to a range of hues, and the responses within a patch are organized by perceptual hue. This is backed up by preliminary 2P results; the two are comparable in spatial scale and color organization (figure 2.23). However, we should bear in mind that these results are not certain. For instance, earlier we saw results about the patchy organization of CR from Landisman and Ts'o (2002). While the OI in that study did not reveal whether individual patches had preferences for single hues or multiple hues, corresponding electrophysiology guided by locations of CR patches from OI indicated that red/green opponent responses occurred in different patches from blue/yellow opponent responses (Landisman and Ts'o, 2002), which differs from the conclusions of the Xiao et al. (2007) and Chatterjee (2010); Chatterjee et al. (2008) studies. This discrepancy has not yet been
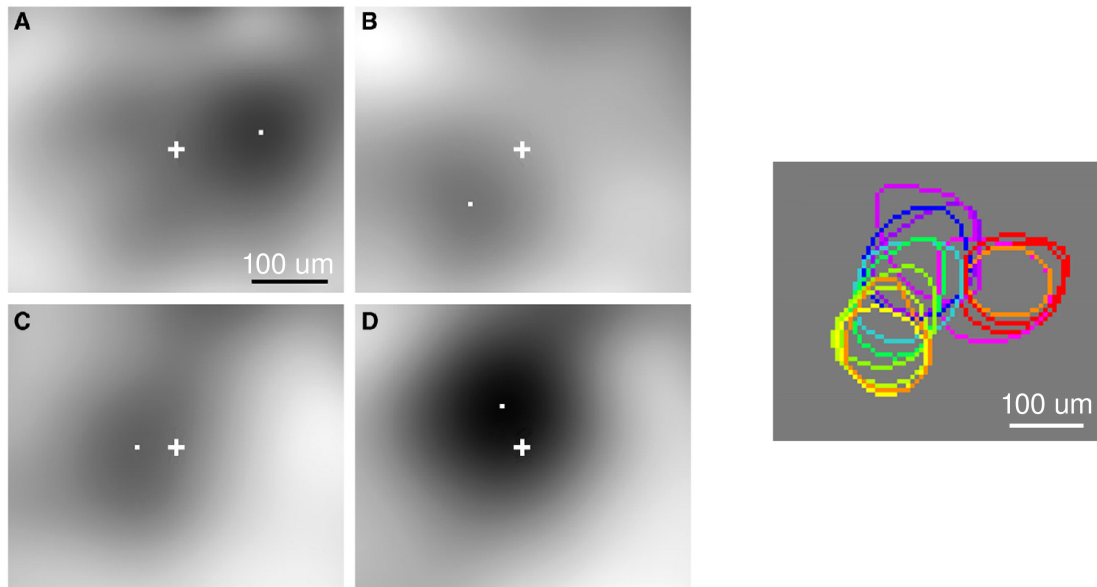
(a)
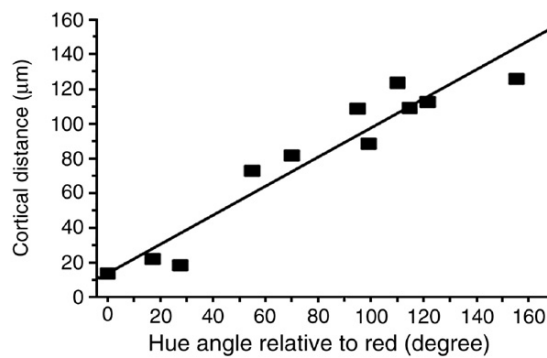
(b)

Figure 2.21: *Organization of hues within CR patches (Xiao et al., 2007). Within color-selective regions, peak responses to different hues are spatially separated, and are arranged according to perceptual hue space (i.e. perceptually similar colors are closer together than perceptually dissimilar colors). (a) Left: Magnified view of top-right white rectangle from 2.18 i.e. one patch, showing V1 responses to uniform stimuli of different colors (A–D: red, yellow, green, blue). Peak is marked with a white square. Scale bar $100\,\mu$m. Right: summary of peak response regions (contours at 75% of peak response for each color) for this patch. (b) Distance between peaks on the cortex correlates with distance between the hues in PCS. Measured across all patches, the correlation between hue angle and cortical distance between peaks indicates that the result illustrated in (a) holds across patches. Images from Xiao et al. (2007).*
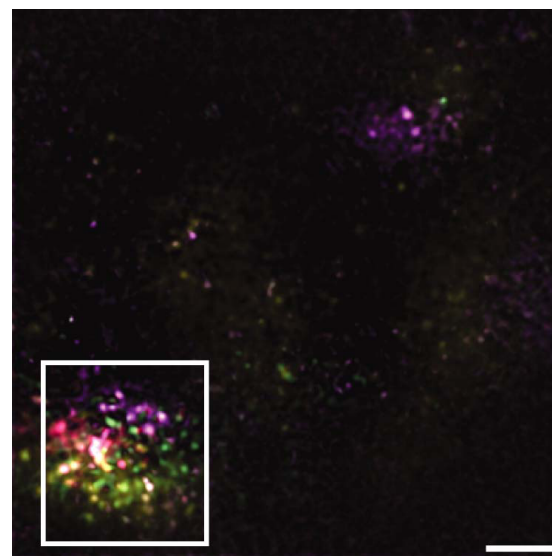
(a)
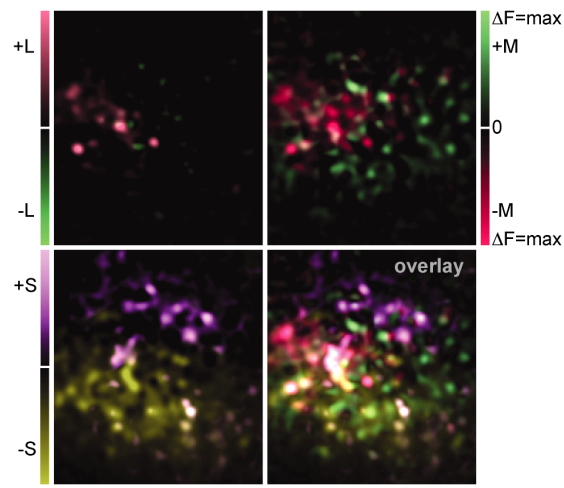
(b)

Figure 2.22: *Organization within CR patches (Chatterjee, 2010). Preliminary 2P imaging study. (a) Combined responses to spatially uniform isoluminant color flashes of various L, M, and S cone contrasts (preceded and followed by gray), showing color responses in patches. (b) Responses to individual stimuli making up the combined response shown in white rectangle of (a). While the patterns were defined by L, M, and S contrasts rather than by perceptual color, the results are still consistent with patches being organized by perceptual hue. The organization of color responses into patches is also consistent with earlier results from intrinsic optical imaging. Scale bars are* $100\,\mu m$*. Figures reproduced from Chatterjee (2010).*

adequately explained, and could represent either genuine differences between animals, sampling issues, or other differences in the methods.

2.6.4 Summary

We can summarize the key experimental results about macaque parafoveal V1 with the following four findings:

ER1-ORmap Spatially contiguous OR preference maps with reasonably uniform selectivity. Widely accepted.

ER2-patches Spatially isolated patches of CR cells (revealing their higher response to chromatic gratings than to achromatic gratings). Reasonably well accepted (a more moderate form of earlier claims that CR and OR are separate in V1).

ER3-range Each CR patch responds to all, or at least a wide range of, perceptual hues. This is a less certain result, derived only from the most recent OI and 2P studies.

ER4-PCS Within CR patches, an organization according to perceptual hue: distance between the peak response to each hue is correlated with perceptual distance. This means that the peak responses to red and violet are closer to each other than either is to green, i.e. implies joining up of opposite ends of wavelength spectrum. Again, as above, this is a less certain result, with support from only a single lab at present.

Having reviewed experimental data relating to the organization of OR and CR in macaque V1, we will now go on to review previous modeling to try to understand this organization. We will be comparing modeling results to these experimental findings.

2.7 Models of CR and OR in V1

In the previous section we saw how OR and CR are represented in adult macaque V1. However, a number of reasons motivate us to consider theoretical approaches to understanding this. Firstly, the existing experimental studies do not address how wiring for this organization of OR and CR could develop—we see only the adult state. We saw earlier that the wiring for color vision develops based on visual experience (section 2.4), so this is an important question. There are experimental results showing the progression of development of orientation, motion direction, and ocular dominance
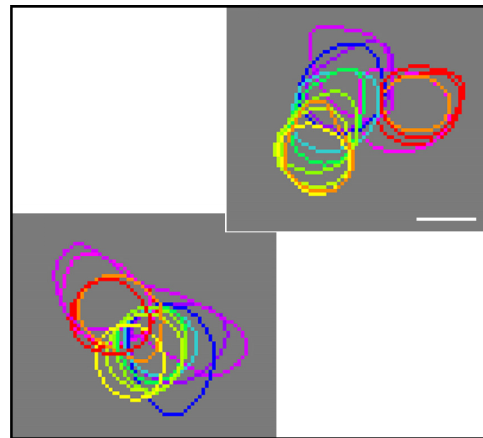
(a)
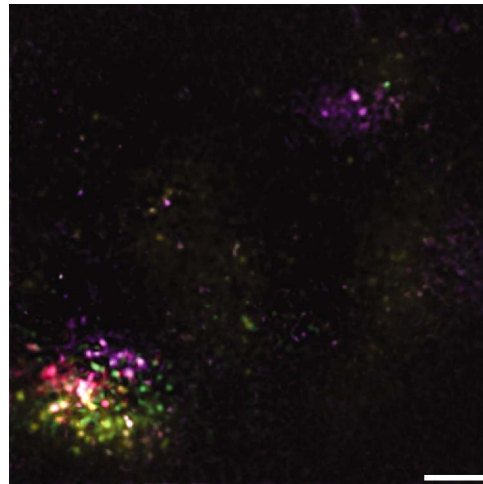
(b)

Figure 2.23: *Comparison of experimental data: CR patch size, spacing, and hue organization. A range of hues is represented in each color patch, and hues are organized according to a perceptual color space. Each 2.5-inch outer square is $3\,\text{mm} \times 3\,\text{mm}$ of macaque V1. (a) iOI responses to spatially uniform hues (contours drawn at 75% of peak), from Xiao et al. (2007). The two gray regions correspond to the small rectangles from figure 2.18. (b) 2P responses to stimuli of varying LMS cone contrasts, from Chatterjee (2010).*

maps over time, before and after eye opening in some cases (Huberman et al., 2008), but the processes underlying such development are only starting to be uncovered. Secondly, we saw that results relating to the organization of CR and OR are uncertain (particularly the representation of hue). Given these, theoretical approaches can be useful to provide insight and potentially guide experimental investigations. Our aim is to use computational modeling to generate hypotheses about how color vision develops.

Previous theoretical approaches to investigating aspects of color vision have included:

- Abstract models based on information processing of physical data (e.g. natural spectra) or simulated cone responses, using techniques such as PCA and ICA (e.g. Buchsbaum and Gottschalk, 1983; Doi et al., 2003; Wachtler et al., 2001).

- Models of adult color processing (e.g. De Valois and De Valois, 1993; Mancuso et al., 2010).

- Models of single neurons or small networks of neurons, often to investigate specific phenomena such as aspects of color constancy (e.g. Courtney et al., 1995; Stanikunas et al., 2004).

- Large-scale neural network models (discussed in detail below).

For our purposes, we dismiss the more abstract models (e.g. independent components analysis (ICA) studies) because while they provide insight into the purpose of V1's processing, they do not provide a mechanism showing what neurons might actually be doing. Single-neuron or small-network models cannot model the kind of measurable, large-scale features such as maps that we are interested in, so we dismiss these for our purposes too. Models of adult color processing cannot explain how the circuitry necessary for color vision actually arises, and furthermore, such models with fixed wiring rule out a large number of visual phenomena that require adaptation, such as color aftereffects. The last category, large-scale neural network models, includes models that have focused on trying to understand how organization for color might arise in the visual cortex based on constraints from physiology and incoming data. We will focus on this category of model.

As we review previous work in this area, the question we will keep in mind is, if we use a biologically plausible, mechanistic model, can we model the emergence of the V1 organization for CR and OR described earlier? Adequately constraining such a model in the absence of, for example, total knowledge of the connectivity, or a complete picture of the developmental processes, poses a challenge. Adding adult behavioral

constraints, in addition to those coming from anatomy and physiology, is therefore desirable. To support this, the model must be able to process realistic input (i.e. natural images)—meaning that it must be able to process features of images other than color alone. Therefore, we restrict our review of previous models to those from this category that work with realistic input, can process OR as well as CR, have a mechanism which could plausibly be implemented in V1, and are of a large enough scale for feature maps. We will now review several such models, with the aim of determining how much they currently explain, and what is the state of the art. Table 2.2 summarizes these models.

2.7.1 CR modeling background

Saarinen and Kohonen (1985) provided the earliest self-organizing model of color, and so was the first step on the path that the models below subsequently followed. Inputs were random artificially generated wavelength spectra (unimodal, bimodal, and polymodal). Three cone types were used, with approximate but not real human cone sensitivities. The model also had RGC, luminance, and chromatic opponent pathways. The model of cortex was small, 15×15 units, and the SOM neighborhood covered the whole cortex. Additionally, there was only input pixel to the cortex, hence it was not retinotopically organized, and there was no possibility of investigating OR. However, this model is interesting because it did show some ordering by dominant wavelength (and also purity) in a self-organized model of cortex for the first time.

Barrow et al. (1996) modeled the development of OR and CR cells in a self-organizing model of V1. Patches of uncalibrated RGB images were presented to a model of the early visual system consisting of red and green photoreceptors (from RGB images—not LMS), center-surround red/green cone opponent and LUM RGC processing, and a model of V1 with adaptable afferent connections, plus fixed, isotropic long-range inhibitory and short-range excitatory lateral connections. The model developed units with OR-selective receptive fields (RFs) that vary smoothly across the cortex, along with patches of unoriented color RFs (each patch containing either red or green RFs).

The Barrow et al. (1996) model successfully showed how the statistics of natural images could affect the RFs for CR and OR. Hebbian learning of afferent connections gives weight patterns related to the principal components of the images: oriented OR RFs and unoriented CR RFs. Lateral excitatory connections cause smooth variation in the RFs, while inhibitory lateral connections prevent all weights from representing the dominant input component. However, the model used RGB input images rather than

	Feature	S85	B96	B05	R05	DP07	R12	E14
1	Image input	✗	✓	✓	✗	✓	✓	✓
2	Natural scenes	✗	✓	?	✗	✓	✓	✗
3	Num. images	-	2	?	-	21	?	112
4	Retinotopic	✗	✓	✓	✓	✓	✓	✓
5	Orientation	✗	✓	✓	✗	✓	✓	✗
6	Num. cone types	3	2	2	3	3	3	3
7	LMS cones	✗	✗	✗	✗	✓	✗	✗
8	Spatial opp. RGC	✗	✓	✓	✓	✓	✗	✗
9	Long-range lat. conns.	✗	✓	✓	✗	✓	✗	✓
10	Adaptable lat. conns.	✗	✗	✓	✗	✓	✗	✗

Table 2.2: *Large-scale models of the development of CR in visual cortex. A number of features are relevant for modeling CR and OR development in macaque V1. (1) If the input contains spatial as well as color (spectral) information (i.e., is in the form of images), then OR as well as CR can be modeled. (2) We consider natural scenes to be realistic input; this is discussed further chapter 3, along with (3) how the number of images affects whether the sample is representative of reality. (4) Retinotopic models have some mapping of spatial location on the photoreceptor array to the cortex, as is found in animals. (5) Some models consider only the development of a representation of color, and not of orientation. (6) All models include at least two cone types, but (7) most do not use realistic LMS cone sensitivities. All the models include some form of cone opponency, but (8) some do not include spatial opponency, a key feature of the RGC for at least luminance processing. (9) Finally, many color (and other visual) phenomena are not restricted to the "classical" receptive field of cortical neurons, and therefore whether the cortical model includes long-range lateral connections (which form the bulk of input to cortical neurons in actual cortex Gilbert et al., 2009), and (10) whether or not those connections can learn, is potentially important. Key: **S85** Saarinen and Kohonen (1985); **B96** Barrow et al. (1996); **B05** Bednar et al. (2005); **R05** Rao et al. (2005); **D07** De Paula (2007); **R12** Rao and Xiao (2012); **E14** Eguchi et al. (2014).*

realistic LMS cone responses. Additionally, the model had only two channels (red and green). CR patches in the model preferred either red or green input (and the balance depended on the images chosen for input), whereas CR patches in macaque V1 respond to many hues. Furthermore, the model's lateral connections were fixed, and hence could not reflect correlations in the input images. Whether e.g. similar colors are preferentially connected during development, or whether there are adaptable patterns in the connection between OR- and CR-selective cells may be important for a number of visual phenomena (e.g. the McCollough Effect; chapter 6).

Bednar et al. (2005) modeled the development of OR and CR maps in the Laterally Interconnected Synergetically Self-Organizing Map model (LISSOM; Miikkulainen et al., 2005). Like the Barrow et al. (1996) model, the Bednar et al. (2005) model was trained using natural images, had two input channels (R and G of RGB), and had cone- and spatial- opponent processing in the RGC. Again, similar to the Barrow et al. (1996) model, this model developed red and green CR patches of unoriented cells in a generally OR-selective, spatially contiguous map. Additionally, unlike the Barrow et al. (1996) model, the Bednar et al. (2005) model showed how the development of lateral connections could lead to preferential lateral connections between neurons of similar CR preference as well as similar OR preference. Understanding such lateral connections is likely to be important for understanding V1 (Olshausen and Field, 2005), and many color phenomena depend on context, suggesting the "extra-classical" receptive field is important. However, as with the Barrow et al. (1996) model, having only two input channels that did not match human/macaque cone sensitivities limited the model's ability to explain results from macaque V1.

Rao et al. (2005) modeled how CR preference (but not OR preference) could develop in a self-organizing model of visual cortex. This model did not use natural input— spatially uniform patterns of random input activation were used. As with the previous two models, the Rao et al. (2005) model model had spatial- and color-opponent RGC channels (a luminance channel, and red/green and blue/red+green channels). After training, an RGB preference map (i.e. not using realistic LMS receptor sensitivities) showed preferences to six discrete RGB values, with similar colors near to each other. However, unlike the Barrow et al. (1996) and Bednar et al. (2005) models, the Rao et al. (2005) model did not develop an OR map, which makes it difficult to compare the spatial organization results to those from macaque. The Rao et al. (2005) model also did not simulate realistic lateral connections, nor did it use a plausible incremental mechanism for learning.

De Paula (2007) extended the Bednar et al. (2005) model to include a trichromatic path-

way with realistic LMS photoreceptors and corresponding RGC pathways. The model used natural images as training input (a selection from the McGill Calibrated Image Database; Olmos and Kingdom, 2004 (MCID)). Like the Barrow et al. (1996), Rao et al. (2005), and Bednar et al. (2005) models, the De Paula (2007) model used cone-opponent center/surround RGC/LGN processing. However, unlike previous models, realistic LMS sensitivities were used. Additionally, this model used two eyes. We will be comparing the results of this model to experimental data in detail in the next section, but briefly, the model developed an OR preference map with patches of CR-selective cells.

Rao and Xiao (2012) also modeled OR and CR maps in a self-organizing model of V1. Training consisted of a pre-natal period during which RGB colored gratings are presented, followed by a visual experience training period during which RGB images of natural scenes are presented. Hence, unlike the De Paula (2007) model, the Rao and Xiao (2012) model used RGB images as input rather than realistic LMS cones. The model contained two cone-opponent channels (red/green and blue/red+green), and a luminance channel. However, unlike the Barrow et al. (1996), Bednar et al. (2005), and De Paula (2007) models, the RGC channels were not spatially opponent. Additionally, the blue/red+green channel was manually connected to four specific cortical regions ("blobs"), while the red/green channel was connected to the remaining ("interblob") regions. Again, this differs from the Barrow et al. (1996) and De Paula (2007) models, which instead showed how unoriented, color-selective regions could arise through self-organization of afferent connections.

Eguchi et al. (2014) provided the first model of CR to include spiking. Images of indoor scenes were used to train the model. Realistic LMS receptors were not used, but correlations were added to simulate L and M cone correlations. The simulated RGC and LGN included cone opponency, but not spatial opponency. The trained model responded to a range of hues, and some clustering of hue responses developed. However, any hue organization present is difficult to evaluate because it was not in the context of an orientation map as in macaques, and it is not clear that any organization of hues according to PCS developed.

Table 2.2 lists the key features of all the models introduced above. Two of those models—the De Paula (2007) and Rao and Xiao (2012) models—have three input channels, can be trained on natural images, and include OR as well as CR. This makes them potential candidates for being able to study macaque V1 organization of CR and OR. Therefore, we will now go on to review the results of these two models in more detail, comparing both closely to experimental data. We will be referring to figures 2.24

and 2.25, which place model results alongside the experimental results shown earlier.

2.7.2   Comparing models to experimental data

The De Paula (2007) model develops an OR map that is generally selective, and comparable to the macaque results we reviewed earlier. Rows 2–4 of figure 2.24 allow the model results to be compared with experiment. The model also develops a patchy representation of CR. The contour plot in row 2 of figure 2.25 shows that responses to full-field hues are patchy. There is no hue CR selectivity map from experimental studies, but the model's CR selectivity map in row 4 of figure 2.25 shows that the contours in row 2 correspond to regions of high CR selectivity. Additionally, regions of high CR selectivity correspond to regions of low OR selectivity. The neurons in CR patches have unoriented connection fields (CFs). Row 2 shows another important result, which is that a single patch can represent more than one hue: we can see that several patches respond to a (limited) range of hues around their preference. Row 5, the hue preference map, shows more clearly that there are neurons with preferences for hues from red to green, and for violet. While there is a limited set of preferences, neurons do respond to hues other than their preferred hue.

The De Paula (2007) model successfully shows that an OR map and CR patches (experimental results **ER1-ORmap** and **ER2-patches**) can develop from images of natural scenes using a plausible model of the early visual system. Additionally, it shows that individual patches can respond to more than one hue. The model also predicts that cells preferring similar hues as well as orientations are found to be preferentially connected by lateral connections during development. Notably, the De Paula (2007) model also includes ocular dominance (OD), although we will not consider that here as we will focus on the organization of hue. For the phenomena we are investigating, there should not be a systematic interaction between hue and OD (i.e. PCS is broadly consistent between eyes), although eventually it will be important to include OD (just as it will be important to include e.g. motion, and other features—see chapter 6).

However, the model does not account for experimental data relating to the organization of hues within patches. The model's color patches each respond to a restricted set of hues (e.g. red to green, or blue). This does not match data showing all or at least a large range of hues can be represented together in a single patch—experimental result **ER3-range**. Furthermore, experimental data indicates that the distance between peak responses is correlated with the perceptual distance between the colors. The De Paula (2007) model's hue contour plot, on the other hand, shows that responses within

blobs are not spatially organized (e.g. if the contour for response to red is concentric with that of the response to yellow, those hues cannot be distinguished via a place code). Therefore, the model also does not match experimental result **ER4-PCS**.

Beyond the results just mentioned appearing not to agree with experimental data, we also find that it is difficult to compare the model and experimental data because several of the model's analyses did not match a particular experimental procedure. For instance, experiments identify CR patches by comparing the responses to BW and RG gratings, while the model reports CR selectivity. For the organization of hue, the spatial organization of hues inside the patches is not comparable to experimental results easily because there is no measure of peak responses, or of distances between peaks.

Like the De Paula (2007) model, the Rao and Xiao (2012) model develops an OR preference map which is generally selective (figure 2.24, column 1, rows 2–4). The OR selectivity is higher in the interblob regions than in the wired blob regions. Figure 2.24 row 5 shows that the model also develops a hue preference map, and that CR selectivity is higher in the blob regions. The hue map contains a range of hues, but it is not clear that all hues are represented, or what their spatial relationship is.

A number of features of the Rao and Xiao (2012) model make it unsuitable for our purposes. We want to explain how cortical wiring might arise, which a model with pre-specified wiring of CR patches will not allow us to do. Additionally, the model uses RGB rather LMS photoreceptors, which means important input channel correlations will not be present. The opponent pathway is not spatially opponent; each retina pixel (R,G,B) is transformed to (R +G, R −G, R +G −B) as input to the cortex. Again, this means missing out on spatial correlations.

As with the De Paula (2007) model, the Rao and Xiao (2012) model is missing some analyses to compare with experimental data. The peak responses to various hues are not analyzed, as they are in experimental studies. Also, the spatial relationship between the hues is not measured.

In summary, the model of Rao and Xiao (2012) matches experimental result **ER1-ORmap**, but for our purposes of explaining experimental findings it does not match **ER2-patches**. And while it appears that each blob represents a range of hues, it is not clear how many or exactly what hues are represented, or that there is a meaningful spatial relationship between the hues. Hence, experimental results **ER2-patches**, **ER3-range**, and **ER4-peaks** are not matched.

Figure 2.24 *(following page)*: *Comparing models to experiment 1: OR vs CR organization. Two previous modeling studies (De Paula, 2007; Rao and Xiao, 2012) have addressed the development of OR and CR, showing OR and CR preference maps while having three input channels and having been trained on natural images, but they have not been compared alongside experimental data as is shown here. The experimental data (right four columns) is reprinted from figure 2.20; each 1.5-inch square is $3\,\mathrm{mm} \times 3\,\mathrm{mm}$ of cortex. Model data (left two columns) is reprinted from the publications indicated in column headings but scaled by approximately matching the OR periodicity. Both models develop an OR selectivity map with regions of low OR selectivity, broadly consistent with experimental data. OR preference maps represent all orientations, and share some features in common with experimental maps.*

	OR sel	OR pref	OR pref & sel
Blasdel (1992b)	NA		
Landisman & Ts'o ('02)		NA	NA
Xiao et al. (2007)	NA	NA	NA
Lu and Roe (2008)			
Rao and Xiao (2012)	NA		
De Paula (2007)	RG-BW / NA		
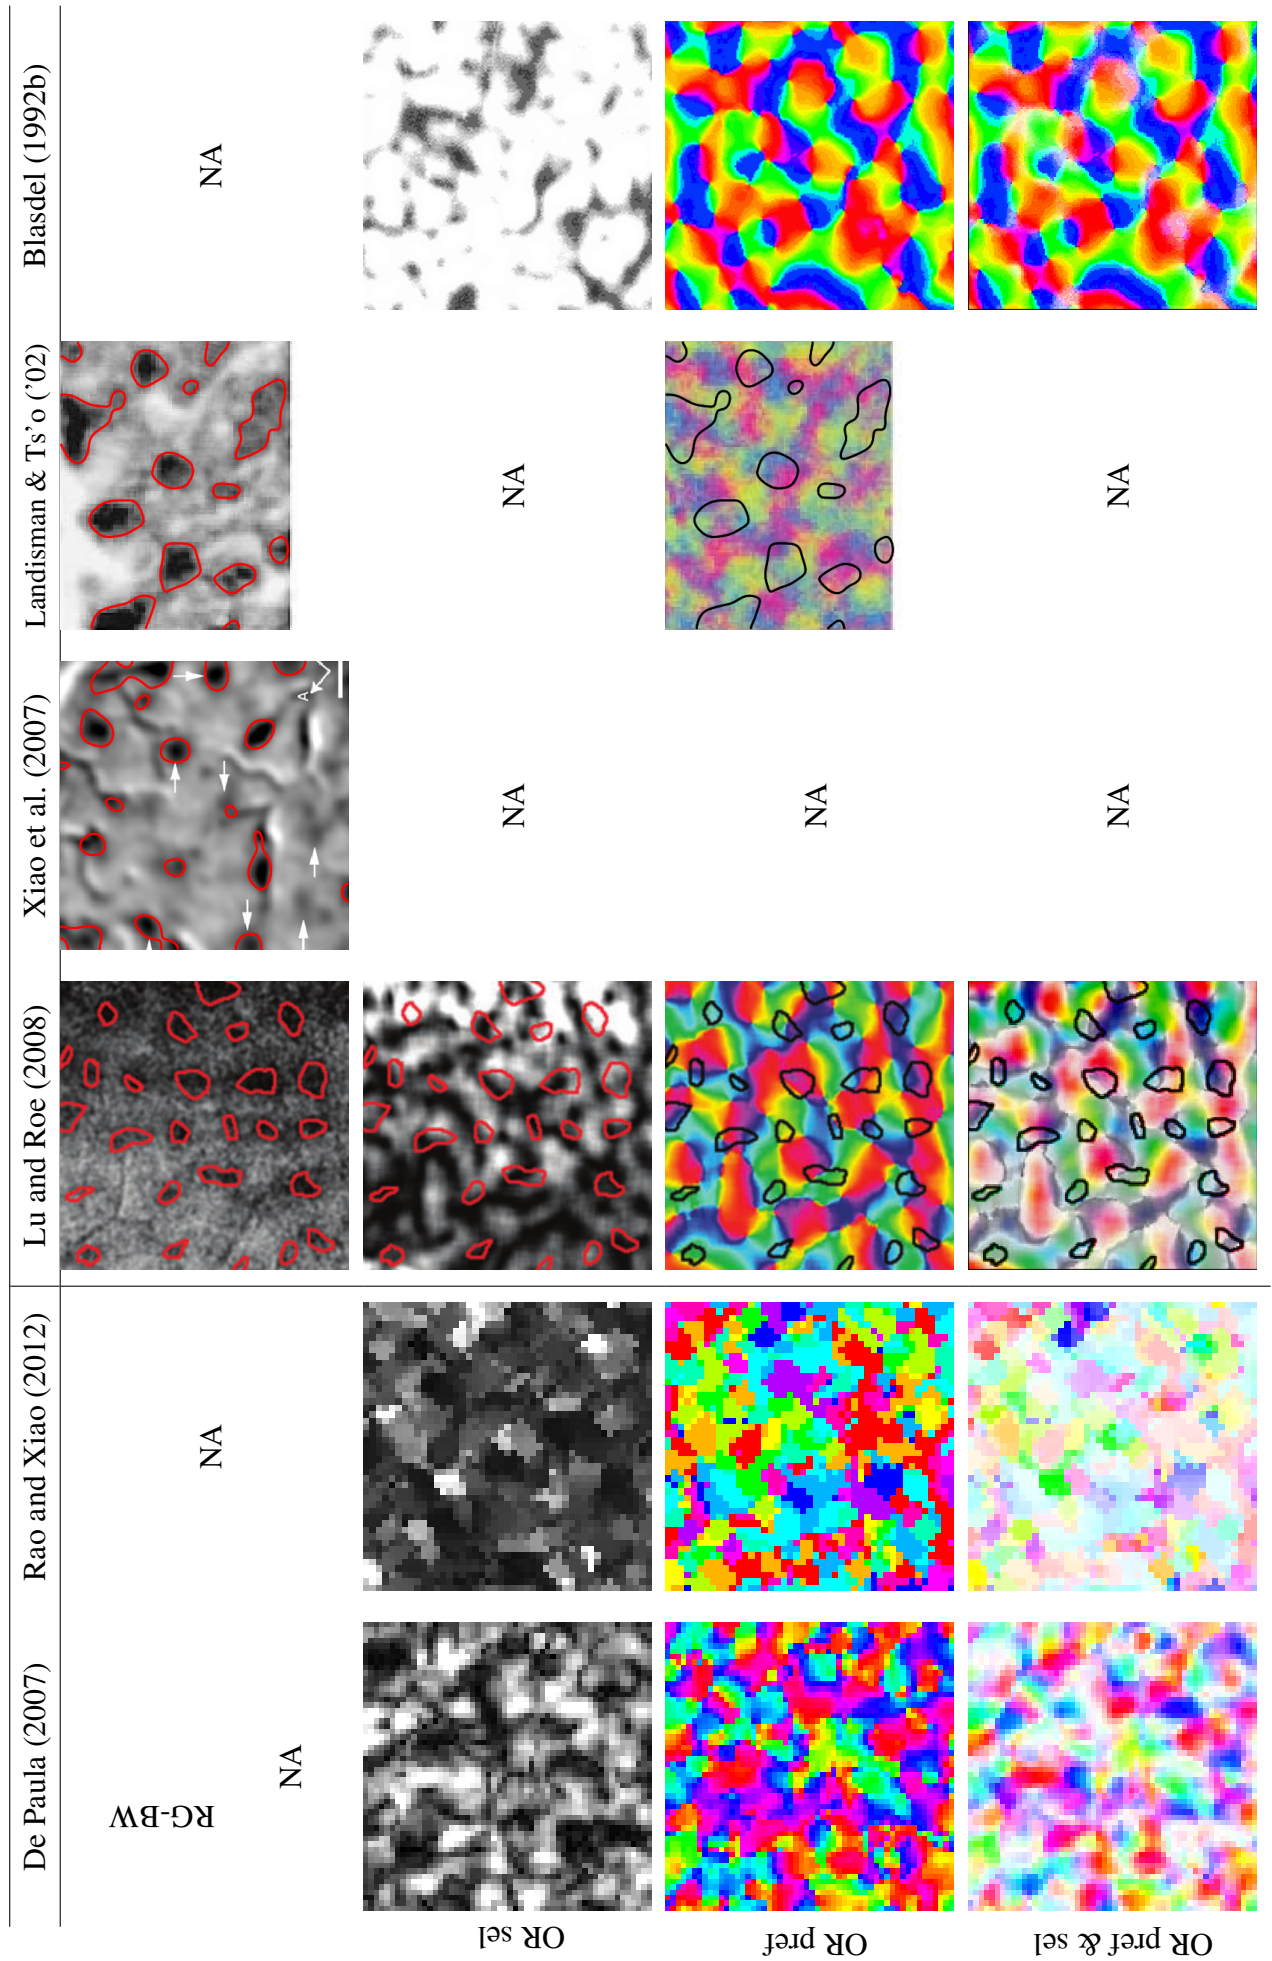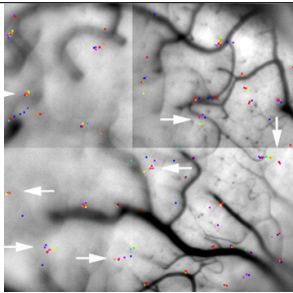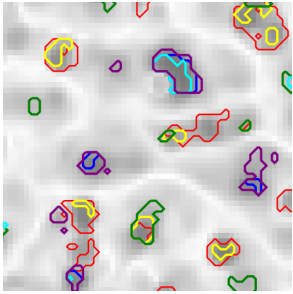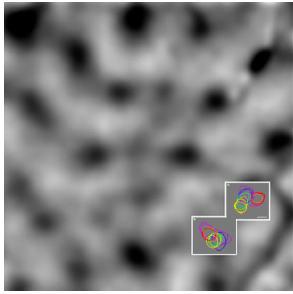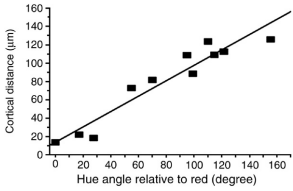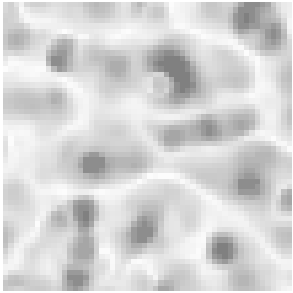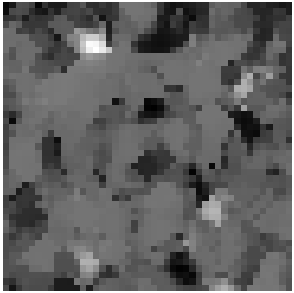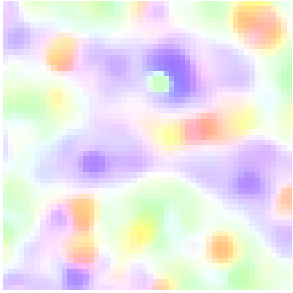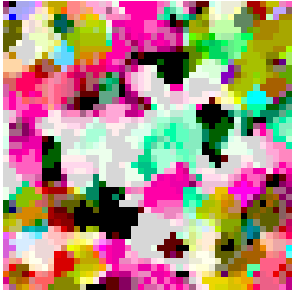
Result	De Paula (2007)	Rao and Xiao (2012)
ER1-ORmap	✓	✓
ER2-patches	✓	✗
ER3-range	✗	✗
ER4-PCS	✗	✗

Table 2.3: *Summary of results from existing trichromatic models of OR and CR map development. The De Paula (2007) model can currently explain how OR preference maps develop, along with patches of highly CR-selective cells. However, each patch does not represent all or many hues, and hues are not organized in a way that corresponds to PCS (as they appear to be in macaque experimental data). The Rao and Xiao (2012) model can explain the development of OR preference maps, but its CR patches are pre-specified, and OR and CR develop in separate layers of the model.*

2.7.3   Summary of modeling status

We have seen that there are a number of models of the development of CR and OR, and studies using these have been successful in explaining several phenomena. However, we have seen that not all the models have the features we need to be able to study OR and CR organization in macaque V1, and none has been closely compared to the experimental data we care about. We narrowed down the selection of models to two, based on their features, and then considered whether they could explain the results we care about. Table 2.3 summarizes our conclusion that the best starting point for our

---

Figure 2.25 *(following page)*: *Comparing models to experiment 2: hue organization. Both models from figure 2.24 (De Paula, 2007; Rao and Xiao, 2012) show some organization for hue, but do not appear to match experimental data, which shows each CR patch representing a wide range of hues. Row 1: peak responses to different full-field hues not reported for models. Row 2: the De Paula (2007) model's response contours to different hues are typically concentric, not offset to allow a spatial code for hue. Row 3: distance between hue peaks (experimental plot introduced in figure 2.21) is not measured for these models, but is unlikely to match the data given hue organization shown in other plots below. Rows 4 and 5: CR selectivity is not available experimentally, except for small samples of cells from electrophysiology. The De Paula (2007) model shows patches of high selectivity. Each patch has one, or a limited range of, preferences. Some preference transitions follow PCS, but preferences for all hues are not present. The Rao and Xiao (2012) model has four hard-wired CR patches (i.e. patches do not arise through activity-driven development). Not all hue preferences appear to be present, but it is not clear how to interpret the hue preference scale. Modeling data reprinted from studies indicated in column headings, scaled by the same factors as for figure 2.24. Experimental data is reprinted from Xiao et al. (2007) (second row contains two images scaled and overlaid).*

	De Paula (2007)	Rao and Xiao (2012)	Xiao et al. (2007)
CR peaks	NA	NA	
CR contours		NA	
CR peak distances	NA	NA	
CR sel			NA
CR pref & sel			NA

work is the De Paula (2007) model.

Despite the successes of the models reviewed above, including the De Paula (2007) model, several problems and unanswered questions remain in modeling the development of CR and OR organization in macaque V1. These fall into three categories:

- *PCS representation problems*: ways in which the models do not match our current knowledge of how CR is represented in macaque V1.

- *Not directly comparable measurements*: ways in which the current analyses do not allow comparison with experimental data.

- *Modeling problems*: ways in which the models do not behave in a reasonable way.

Each group contains a number of problems, which we now list, starting with differences between PCS representation in the models vs. in macaque:

P1-range The De Paula (2007) model is the only one to develop CR patches alongside an OR map while also meeting other key feature requirements such as being trichromatic, and using natural image input. However, even in the De Paula (2007) model, each CR patch only responds to a limited range of hues around its preference (e.g. red and yellow, or blue and violet). This contrasts with **ER3-range**, which shows patches responding to many/all hues. None of the other models that develop CR patches shows a hue range within the patches matching experimental data.

P2-PCS None of the models appears to have hue spatially organized according to PCS (**ER4-PCS**). For instance, the De Paula (2007) model's contour plot (figure 2.25) shows largely concentric responses to different hues. There are some apparently-smooth-in-PCS transitions in the CR preference map (figure 2.25), but it is not clear that they cover PCS or that across V1 the distances between colors correlate with PCS. We saw that Eguchi et al. (2014) had some organization for hue, but it is not clear whether it follows PCS, and in any case this organization is not in the context of an OR map (**ER1-ORmap**), or discrete patches (**ER2-patches**).

P3-pref If the coding of hue is to be useful for PCS, we might expect a hue preference map that shows preferences for all hues, similar to maps for OR preference. Of course, the preferences do not have to be uniform, but may be biased e.g. in a similar way to OR preference maps containing more vertical and horizontal preferences than oblique preferences (Chapman and Bonhoeffer, 1998). The hue preference map from Rao and Xiao (2012) appears to be missing a range of hues,

though it is not clear how to interpret the hue map's scale. The De Paula (2007) model also has preferences for only some hues, suggesting both models develop a biased representation of hue. Interestingly, however, De Paula (2007) found that a preference map containing neurons with preferences for all hues could develop under specific circumstances. A model trained on 21 images with a modified distribution of hues (distribution unknown), developed a hue preference map containing preferences for all hues, although individual CR patches still preferred a limited range of hues (figure 2.26). The altered hue distribution, although unknown, is unlikely to reflect biological reality because it was created from a small set of images by randomly rotating each image's hue once. Also, no corresponding OR map was shown.

While it appears that previous modeling work does not match important results about organization for PCS in V1, there are difficulties comparing the models with experimental data because the analyses do not correspond. For instance, we have been making statements about the models' representation of hue based on hue preference maps, but there are currently no experimental hue preference maps. These problems are grouped together as *not directly comparable measurements*:

P4-peaks No analysis of location of peak response to each hue (and distances between peaks). Missing from all previous models.

P5-RGBWgratings No analysis using RG vs. BW gratings to identify color regions, as in experimental work. Missing from all previous models.

P6-HSV All analyses were done in HSV space rather than using a perceptually uniform color space, or at least matching specific color stimuli used in particular studies. HSV (a) does not properly separate perceptual attributes of color, (b) is not perceptually uniform, and (c) is a device-dependent space (usually it is not specified in relation to a standard, absolute RGB space such as sRGB).

Finally, there are also a number of problems from a modeling point of view, *modeling problems*:

P7-LMSscaling The De Paula (2007) model requires ad hoc scaling of the LMS channels. These are set manually to achieve a CR preference map in which each cone type is represented, and the settings apply only to the specific (small) set of images that was used. The same problem exists in Bednar et al. (2005), and likely also exists in Barrow et al. (1996) because the authors briefly mention an unbalanced number of red and green patches, and used only two images. We do not know if this problem exists in Eguchi et al. (2014) or Rao and Xiao (2012), but

(a)

rotate hue

(b)
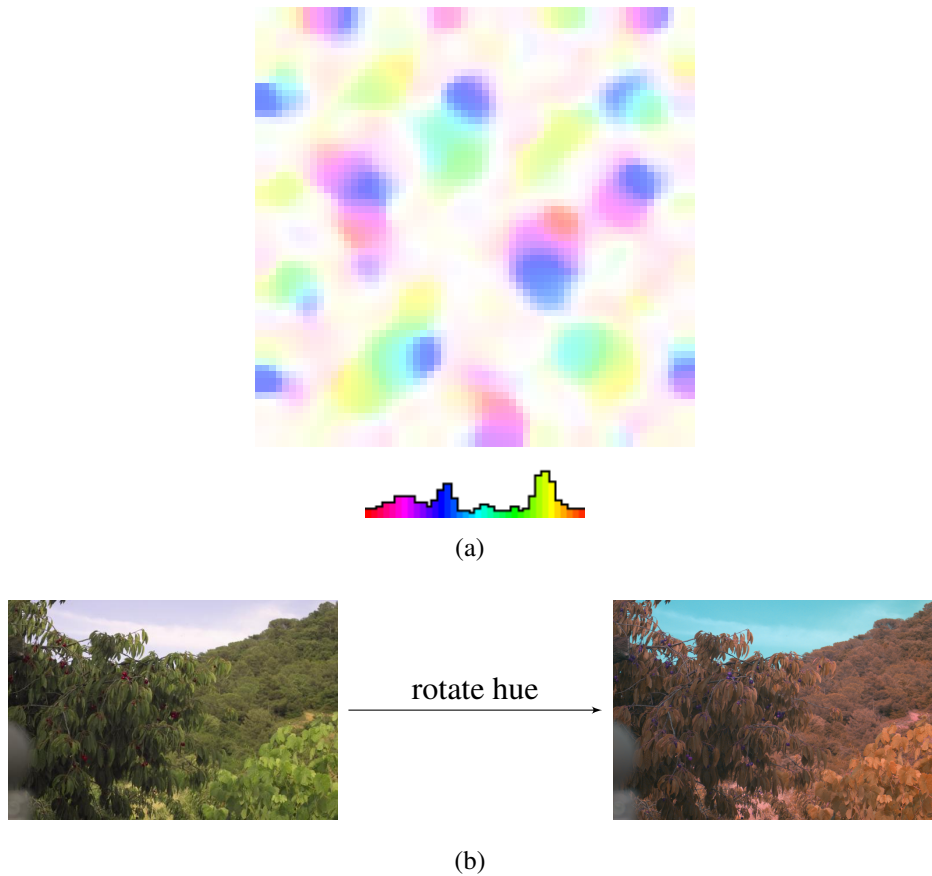
Figure 2.26: *Effect of input hue distribution on input-driven models. De Paula (2007) shows preferences for all hues can develop by manipulating the input hue distribution. (a) The LISSOM model can develop preferences to all HSV hues: a hue preference and selectivity map is shown with a histogram of preferences beneath (reprinted from De Paula 2007). CR patches still do not each respond to all or many hues, but do at least respond to a greater range. No corresponding OR map was shown. As illustrated in (b), for this simulation the hue distribution of a set of 21 images was altered by rotating the hue of each image once. That is, each RGB input image (example from MCID on the left) was converted to HSV, and H, which varies circularly between 0 and 1, had a single uniform random number between 0 and 1 added to it, and then the image was converted back to RGB (on the right). It is not clear what hue distribution this resulted in, although it is unlikely to have been either uniform in hue space, or representative of an individual's visual experience. We will explore the effect of input hue distribution later in this thesis.*

our work in the next chapter indicates it is highly likely to be a general problem.

P8-V1StableDev While the developmental path of CR maps is not known, organization of other feature preferences is known to develop in a stable way in animals (e.g. orientation: Chapman et al., 1996, motion direction: Li et al., 2008). Therefore, it seems reasonable to assume the same for CR. However, the LISSOM model of V1 used by De Paula (2007) does not develop in a stable way (Stevens et al., 2013b). The correct activation level of V1 in LISSOM is maintained by periodic manual parameter adjustments, and the OR and CR map states can change suddenly at these adjustment points. This makes it difficult to investigate the development of hue representation, because the hue representation can change for spurious reasons. We do not know if this problem applies to other models, but e.g. the Rao and Xiao (2012) model uses a SOM model with shrinking lateral radius and no homeostasis, and so seems likely to suffer the same problem (though perhaps more gradually, rather than at discrete times).

Solving these *modeling problems* is important for three reasons. The first is that they make investigating a model difficult (e.g. one cannot easily vary input images, because several other parameters must be set manually via an iterative, ad hoc procedure to match). The second reason is that eventually we want to combine the model of CR and OR with a general-purpose model of the cortex incorporating other dimensions (Bednar, 2012), and it is infeasible to manually set a variety of parameters for each dimension in concert. The third reason is that reducing the number of parameters (e.g. by implementing neural homeostasis to avoid having to manually control activation levels) reduces the modeler's degrees of freedom, while making the models simpler and more realistic.

In this thesis, we will go onto address these problems, starting from the De Paula (2007) model. Given the overall realistic architecture, and the results that are beginning to match the biological data, extending this model seems like the best approach. In chapter 4, we will fully describe and then extend the De Paula (2007) model to begin addressing the *modeling problems*, which involves simulating adaptive homeostasis in the input layers and V1. We will also add analyses matching experimental procedures. However, before this, we must investigate what input to use to simulate the visual experience of an individual, which will be the topic of the next chapter.

2.8 Conclusion

In this chapter we have seen a color vision pathway that starts from light, which is captured and transformed by the early visual system into our experience of color, PCS. We also saw that the development of color vision depends critically on both the specific visual input and physiology—and yet despite large varieties in both, individuals develop a reasonably consistent PCS. Organization for PCS appears to arise in the brain as early as V1; experimental studies of macaque V1 find CR cells in patches alongside a spatially contiguous organization of OR. The representation of hue within those patches appears to be organized by PCS, but these results are still debated. Meanwhile, computational modeling of the development of V1 indicates how CR patches can develop alongside OR maps, but not how hue maps could form within CR patches. Along the way, we have found a number of puzzles to address. Characterizing the color statistics of visual experience is investigated in the following chapter, and then extending and analyzing the most successful computational model so far in order to understand where hue maps come from begins in chapter 4.

Chapter 3

Colors in natural image databases

3.1 Introduction

Chapter 2 reviewed experimental evidence indicating that color representation in macaque primary visual cortex is organized according to perceptual color space, and that visual input is critical for the development of color vision. The goal of this thesis is to model how such a perceptual organization for color could arise in primary visual cortex, alongside orientation selectivity, through development driven by visual input. To do this, we need to find practical options for simulating visual input to the model. We also need to understand this input in terms of its color statistics, in order to understand how the model will be affected. These are the two main goals of this chapter.

We will begin by reviewing previous investigations of the color statistics of the environment. Since we are focusing on experiments on macaque, this environment could be the evolutionary niche of macaque, or the specific environment of individual macaques used in the imaging experiments introduced in the previous chapter. Unfortunately, imaging studies using macaque do not specify the rearing environment of the individuals, and the authors themselves may not be able to find out (Y. Xiao, personal communication). Animals could presumably be wild caught, raised in a large animal facility, or raised in a scientific laboratory. In terms of visual history, and hence distribution of hue percepts experienced, those environments are likely to be very different from each other, with perhaps the only commonality being seeing other macaques. Previous work characterizing the statistics of images in relation to human vision has generally focused on locations where the visual system is thought to have evolved (Webster and Mollon, 1997), and for lack of any more-specific data we too will assume natural scenes are a reasonable model of the individuals' visual history for our modeling work.

After reviewing previous analyses of natural scenes, we will then investigate what practical options are available for input to the model. As discussed in the previous chapter, input to the visual system is light: a spectral power distribution, varying over time, filtered through the eye's media, absorbed with some probability by one of three cone photoreceptor types. Since we want to model the development of color (CR) and orientation (OR) organization in the brain, the most complete input to a model would be this filtered spectral power distribution (SPD) falling on each photoreceptor, recorded over an individual's lifetime. However, recording such data is impractical for many reasons, so we must find an approximation. In this chapter, we will therefore review alternative methods for obtaining realistic input. We will see that images of natural scenes captured using a calibrated RGB camera, and transformed to simulate LMS cone excitations (proportional to absorptions), are currently the most practical input, giving the best balance of accuracy of simulation of LMS excitations while capturing a large enough sample of the environment.

Having selected an option for input, we will also characterize a number of databases of such input in terms of the statistics of cone excitations and hue percepts they generate (the latter is what we mean by "the colors present"). Understanding the input data in these two ways is critical to understand the development of models of the kind reviewed in the previous chapter. Those models have cortical learning algorithms that are affected by the statistics of cone responses (such as correlation between different cone types). Additionally, since we are evaluating a model's ability to develop a representation of perceptual color space (PCS), we need to know the frequencies of hue percepts generated by these natural image databases. An image database could overall generate hue percepts that are a reasonably uniform sampling of perceptual color space, or could generate hue percepts that overall do not sample perceptual color space uniformly.

In agreement with previous work, we find that the databases have cone excitations that are highly correlated between types, with some cone types absorbing more energy over the entire database, and excitations skewed towards low values. Additionally, we find the databases generate hue percepts that are a highly non-uniform sampling of PCS. We believe this reflects the genuinely "color biased" visual experience of individuals (a bias that will vary depending on environment, and that will vary over time). The bias has important implications for a model that attempts to learn a PCS based on the incoming statistics. In the previous chapter, we saw that the visual environment can critically affect color vision, yet PCS is broadly consistent across individuals from environments with quite different color statistics. Any developmental model will therefore have to do an important job: be able to organize a reasonably consistent PCS

from variously biased input datasets.

3.1.1 Contributions of this chapter

- Section 3.2 summarizes previous investigations of the statistics of LMS cone excitations in response to natural scenes, and the distributions of hue percepts resulting from those scenes. Cone excitations are highly skewed towards relatively low values, and are highly correlated. The distribution of hue percepts generated is biased towards blue/yellow, and restricted (i.e. some colors almost never appear). Apart from being a comprehensive review of previous studies of perceptual colors in natural scenes, we also find there is not yet an analysis in the form required to understand the impact of the hue distribution on a model of the development of CR organization.

- Section 3.3 reviews available databases of natural visual input, including calibrated trichromatic images and hyperspectral images of natural scenes. We conclude that calibrated trichromatic images of natural scenes are currently the most practical input to use for a model of the development of OR and CR organization, giving the best balance of accuracy of LMS simulation and wide sampling of scenes.

- Section 3.5: Characterization of the McGill (Olmos and Kingdom, 2004) and Barcelona (Párraga et al., 2010) trichromatic calibrated color natural image databases in terms of cone excitations and perceptual hue distributions. As found in previous studies of natural scenes, cone excitations are skewed towards low values, and are highly correlated. We also find total excitations are unequal between the cone types over the entire datasets. Furthermore, the distribution of hue percepts generated is a highly non-uniform sampling of PCS, being biased towards blue/yellow, and covering a restricted fraction of PCS. We show this using CIE L* C* h_{ab} color space (CIELCh$_{ab}$), measuring the distribution in a form that will help predict its impact on an input-driven developmental model.

- Finally, we also create a database of color-manipulated images that does uniformly sample PCS, but which retains cone correlations, which may be useful as training data for developmental models. We additionally show that using RGB images as input could separate the effect of cone correlations from the effect of perceptual hue distribution.

3.2   Background 1: Color statistics of natural scenes

As reviewed in the previous chapter, LMS cone excitations are the information that the visual system starts with. The next stage in the visual pathway's processing of color appears to involve cone opponency. After cone opponency, it becomes less clear how the visual system is organized for color, with a variety of different possibilities at various stages of processing. By the time percepts of color reach conscious awareness, we can define perceptual color spaces that approximately describe our experience of color. In this section, we will review previous investigations of color representation at these three levels (cone responses, cone-opponent responses, and in perceptual color spaces); characterization of any intermediate stages there may be is so far lacking.

Before we proceed to review previous studies of color statistics in the spaces above, we note that an alternative, complementary approach has been to find spaces based on physical data rather than on physiology or perception. For instance, natural re-flectance spectra have been characterized in various ways (e.g. using principal compo-nents analysis (PCA) or independent components analysis (ICA)) to determine, from an information-processing perspective, optimal ways in which they may be repre-sented. Such analyses can be compared e.g. to known physiological representations, such as having three cone types with highly overlapping spectra, or cone-opponent processing, or the types of cells found in primary visual cortex (V1), and thereby help to explain the purpose of particular processing steps. However, we are not reviewing such work here, instead focusing solely on the physiological and perceptual represen-tation of color. We want to learn about the representation of color in ways that will directly relate to our subsequent modeling work, in order to understand the model.

3.2.1   Recording natural scenes

To determine the color statistics of natural scenes, it is first necessary to record them in some way. There are a number of possible methods:

- Observations of natural scenes can be made by humans. For instance, colors in a scene can be compared by eye to samples of known characteristics (e.g. known reflectance spectrum, or known chromaticity). Such observations are very time consuming, and difficult to do with spatial precision.

- Radiances of natural scenes (and natural illuminants, e.g. different phases of daylight) can be recorded with a spectroradiometer. While allowing accurate re-construction of spectral power distributions (and hence of LMS cone excitations,

via the known cone fundamentals), the result is a coarse sampling of the scene (i.e. very low spatial resolution).

- Reflectances of the contents of natural scenes can be determined using a spectrophotometer. As mentioned in the previous chapters, the human visual system is more concerned with reflectances than with radiances (i.e. we "discount the illuminant" to some extent), so having recordings of reflectances as well as radiances is useful for understanding vision. As with the above technique, though, the result is a very coarse sampling of the scene.

- A hyperspectral camera can be used to capture the radiance of a natural scene in multiple small-wavelength intervals. Hyperspectral imaging is time consuming— for each scene, many images must be acquired while making sure there is no movement (or compensating for it)—so there are currently not many hyperspectral images of outdoor natural scenes. However, the advantage of hyperspectral images is the combination of spatial data with accurate SPDs, which allows accurate conversion into any color space, including LMS cone excitations or CIE 1931 XYZ color space (CIEXYZ). Including objects of known reflectances in the scene allows the illuminant to be inferred, and hence the scene's reflectances can be recovered.

- A trichromatic device such as an RGB camera or imaging colorimeter may be used. The difference from a hyperspectral camera is that instead of multiple narrow wavelength ranges being sampled, only three broad ranges are sampled. An imaging colorimeter captures radiances using functions that match CIEXYZ (and hence its output can be linearly transformed to LMS cone fundamentals). If the spectral sensitivities of a camera are known (and the camera's non-linear responses can be converted to linear ones), these can be approximately mapped to LMS cone sensitivities (e.g. using a linear tranform minimizing the error between the transformed and LMS sensitivities, or by recording camera responses for multiple known spectra and creating a mapping). RGB cameras have the advantage of being readily available, and of acquiring high-resolution images very quickly.

3.2.2 Cone statistics

In the previous chapter, we saw that daylight color vision begins with absorption of photons by L, M, and S cones. Cone excitation is proportional to absorption; after this, the cone may have a nonlinear response. Our model of color vision in the following

Data set	Analysis in	Correlation		
		L/M	L/S	M/S
White noise	Doi (2003)	0.916	0.0726	0.132
DDB	Doi (2003)	0.994	0.828	0.862
CCHID	Ruderman et al. (1998)	0.990	0.565	0.650
BHID	Doi (2003)	0.990	0.853	0.912
MCID Foliage	De Paula (2007)	0.987	0.906	0.931
MCID Fruit	De Paula (2007)	0.978	0.877	0.910
MCID LandWater	De Paula (2007)	0.992	0.901	0.937
MCID Winter	De Paula (2007)	0.999	0.995	0.996

Table 3.1: *Cone excitations are highly correlated. Each row shows correlations between pairs of cone types for different data sets (abbreviations expanded in a subsequent section). The first row shows correlations for white noise (equal energy at every wavelength), illustrating the instrinsic cone sensitivity correlations. L and M cone sensitivities are highly overlapping, hence their excitations are highly correlated. Subsequent rows are for excitations in response to natural scenes (the image databases themselves are discussed in a subsequent section; see table 3.2 for details). Natural images further increase the correlations. L/M correlation is always highest, but there are strong correlations between all channels. Scenes that do not appear colorful (e.g. winter scenes dominated by snow—last row) have particularly high correlations.*

chapter will begin with cone excitations. Previous work simulating cone excitations for natural scenes has found excitations are highly correlated (table 3.1). Given the highly overlapping cone spectral sensitivities (figure 2.10 in the previous chapter), this correlation is not surprising. However, we can see that natural images further increase the correlations, particularly for L and S cones, and for M and S cones: generally, natural spectra contain energy across wide parts or all of the visible spectrum.

In addition to highly correlated cone excitations, previous studies have shown the distribution of excitations is highly skewed towards low values (Ruderman et al. 1998; Garrigan et al. 2010; figure 3.1). We will discuss in later chapters how subsequent coding may alter this distribution, but it is a feature of the input that any computational model will have to deal with.

In the previous chapter, we learned that cone excitations are difficult to relate to perceptual color space. However, understanding their statistics is important for understanding how they will affect a learning algorithm such as Hebbian learning. Now we go on to review previous analyses of natural scenes in spaces closer to our perception: cone-opponent space, and models of perceptual color space.
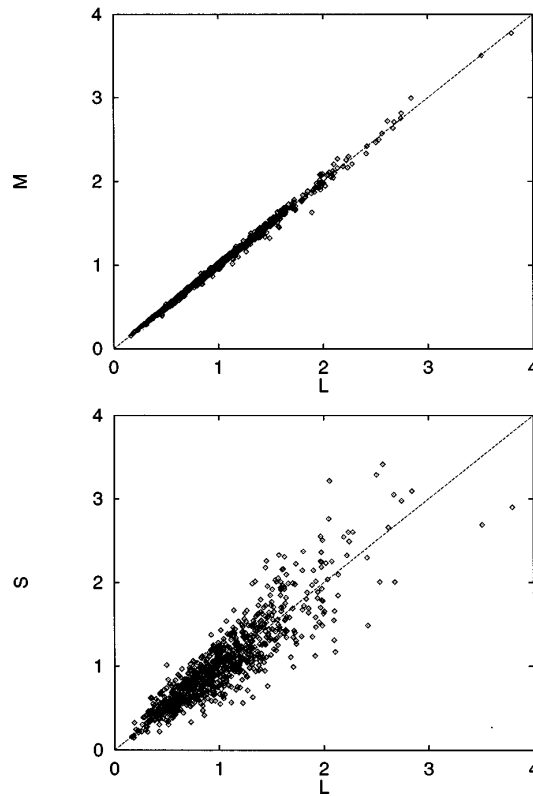
Figure 3.1: *Skewed distributions of cone excitations in response to natural scenes. Analysis of cone excitations for the central* 128×128 *pixels of 12 natural scenes, imaged with a hyperspectral camera system, indicates that cone excitations, apart from being highly correlated between cone types, are also skewed towards low values. The mean of each cone type's excitation is normalized to 1.0, so the mean for each plot is (1,1). Uncorrelated distributions would have a high variance around the diagonal, and unskewed distributions would be distributed evenly about the mean, along the diagonal (contrary to the observed data). Images reprinted from Ruderman et al. (1998).*

3.2.3   Color distributions

We saw in the previous chapter that the visual environment is critical for the development of color vision. For instance, monkeys raised in an abnormal color environment make abnormal judgments of color similarity. In the previous section, we reviewed the statistics of cone excitations for natural scenes. In this section, we will review previous work about the properties of natural images in spaces closer to our perception—in particular about the distribution of hue percepts elicited by natural scenes. However, first we explain what we mean by the hue percepts elicited by natural scenes, and their distribution.

There are various models of adult perceptual color space, which can be used to take e.g. cone excitations or CIEXYZ values and predict hue percepts (chapter 2). Color appearance models are an active area of research, with the most complex for calculating the color at a particular location taking into account features including the background, the surround, stimulation history of the observer, and other features that affect perceived color (Fairchild, 2005). However, studies estimating colors in natural scenes in bulk have used simpler spaces, such three-dimensional cone excitation space (with L, M, and S axes—as reviewed in the previous section), or a cone-opponent space computed from these. Alternatively, studies have illustrated gamuts on chromaticity diagrams (chapter 2), e.g. the *(x,y)* diagram. Finally, some studies have used uniform color spaces such as CIE 1976 L* u* v* color space (CIELUV) or CIE 1976 L* a* b* color space (CIELAB). Spaces like CIELAB and CIELUV, designed to be perceptually uniform and to allow predictions of whether differences between spatially uniform stimuli are distinguishable under specific conditions, give the best approximation of human hue percepts among the aforementioned simple spaces. While not as accurate as the more advanced color models, CIELAB and CIELUV are simpler, and faster to compute (each pixel is calculated independently).

Previous work investigating the colors present in natural scenes has found natural scenes tend to be dominated by only a few colors. To do this, studies have considered both scene reflectances and illuminants. The visual system receives radiances as input (as sampled by the cones), but it is generally accepted that the visual system rapidly adapts to the illuminant, so we are therefore we are interested in both.

Hendley and Hecht (1949) visually compared the colors in natural scenes to Munsell chips of known chromaticity. Scenes were from around New York state, in the summer and autumn, and most were of terrains viewed at a distance (from 0.6 m to 32 km). One color-normal viewer made 1000 observations of 235 objects. Figure 3.2(a) shows that
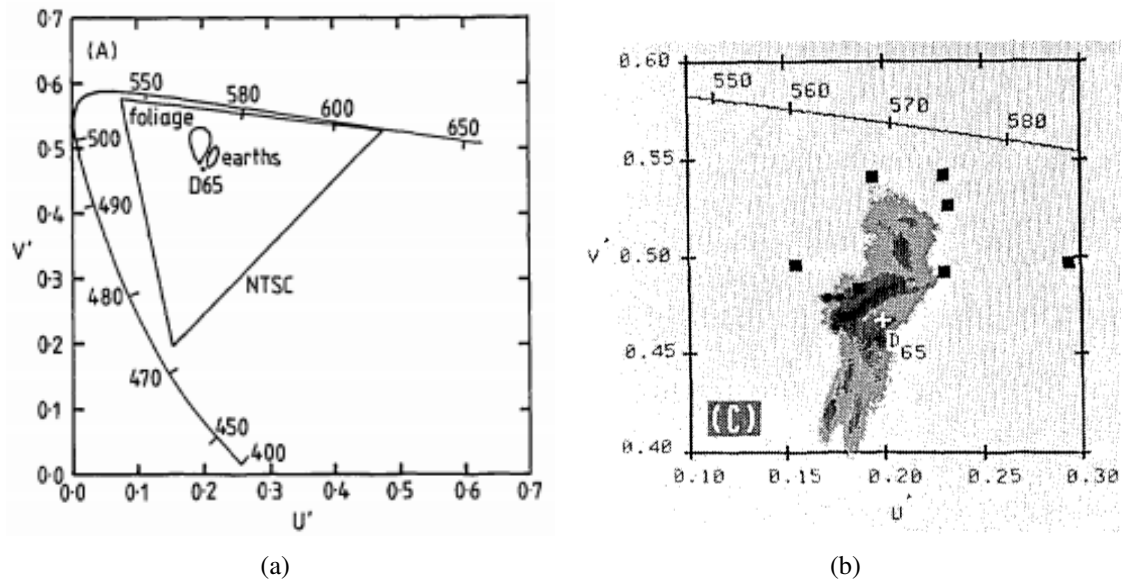
Figure 3.2: *Colors in natural scenes: Hendley and Hecht (1949) and Burton and Moorhead (1987). (a) Gamut of visually observed natural scenes (foliage and earth) shown on the UCS diagram. Also included is the gamut of a television color standard (NTSC), to emphasize the restricted extent of natural colors commonly found within this gamut. Data originally from Hendley and Hecht (1949), but replotted by Burton and Moorhead (1987) (from which image is reprinted).*

(b) 19 images recorded with a film camera system calibrated to CIEXYZ, plotted as a frequency histogram on the UCS diagram as in (a), but a narrower range. Light gray: 100% of gamut, gray: 75% of gamut; black: 25% of the gamut. The images were of wooded and grassland scenes, from a distance of approximately $1 - 5$ km. Again, only a restricted range of all possibly perceptible chromaticities is occupied by the scenes. Image reprinted from Burton and Moorhead (1987).

the gamut of observed colors covers only a small part of the chromaticity diagram (low chroma—particularly at greater distances—and a restricted range of hues). The authors summarize the colors as: yellow-green (vegetation), yellow—orange/red (earths and dried vegetation), and blue (water, sky, distant objects). In autumn, vegetation colors cover an increased range (yellow-green through orange and further into red).

Nickerson et al. (1945) analyzed at close range foliage and earths specially selected by soil scientists to represent the range of soil colors possible. Their results agree with those for earths found by Hendley and Hecht (1949), but additionally include samples with higher chroma, and more red hues are present.

Burton and Moorhead (1987) photographed 19 long-range views of natural terrains (from at least 0.5 km away), and converted three-channel calibrated scans of film to CIEXYZ. Results were similar to those of Hendley and Hecht (1949), as shown in figure 3.2(b): hue variation is restricted and mainly along a bluish-yellowish axis, and chroma decreases with distance.

Webster and Mollon (1997) examined terrains and vegetation in arid mountain and desert scenes, and in temperate rainforest. Viewing distances ranged from meters to kilometers, and generally included sky, foliage, and earth. 19 images from a hyperspectral camera system, and 87 scene measurements with a spectroradiometer, were analyzed. Figure 3.3 shows a similar distribution of color was found to previous studies: restricted, and strongly biased along a bluish to yellowish-greenish axis.

Linhares et al. (2008) analyzed 50 natural and rural scenes captured with a hyperspectral camera and determined the distribution of discernable colors. Images were analyzed in CIELAB, and the CIEDE2000 color-difference formula was used to determine which colors would be distinguishable to a human observer. Most colors appeared only rarely, with the fraction of discernable colors as a function of occurrence being a negative power distribution (Nascimento et al., 2008). Consistent with previous work, the authors conclude that most natural chromaticities are distributed in a small gamut, and that many discernable colors do not appear in nature. However, the authors point out that their images may be missing saturated colors found in e.g. jungle scenes. Also, subsequent work has pointed out difficulties in determining the number of discernable colors in a scene: doing so depends on more than only just-noticeable differences (jnds) in a simple space like CIELAB (see e.g. Marín-Franch and Foster 2010; Masaoka et al. 2013). Despite this, the general conclusion of a restricted distribution of chromaticities, and natural scenes not (or extremely rarely) containing certain colors, still agrees with previous work.
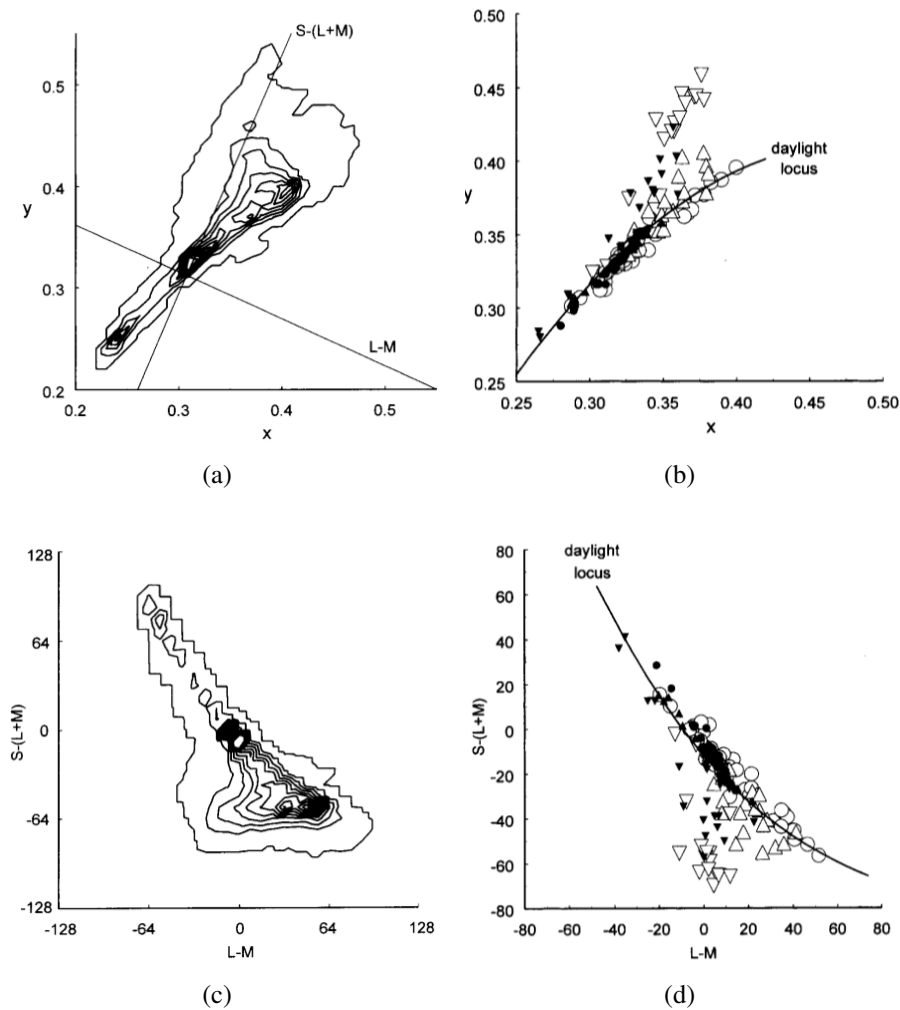
Figure 3.3: *Color in natural scenes: Webster and Mollon (1997). First row: (a) gamut of chromaticities (density contours represent increments of 50) and (b) average chromaticity for ≈ 100 scenes, on the* (x,y) *diagram. As with previous work, the chromaticities in (a) occupy a restricted area of the diagram, and are biased along a bluish/yellowish axis, consistent with sky, foliage, and earth. Average chromaticities in (b) are plotted for different locations with unfilled markers, distinguished by symbol shape, and the chromaticities of illuminants are plotted with filled markers (of the corresponding type). In general, average chromaticity follows the same bias as in previous work (bluish/yellowish), but in lush rainforest (downwards triangles) it is biased towards green. The average chromaticities of illuminants has a smaller range than for the scenes, and follows the daylight locus (chromaticities of phases of daylight), although again, in lush forest it is biased to green. Second row: (c) gamut of chromaticities and (d) average value for each scene on cone-opponent axes (scaled to roughly equate visual sensitivity to each axis). The distribution of responses on cone-opponent axes is again restricted and biased. Images reprinted from Webster and Mollon (1997).*

Yendrikhovskij (2001a,b) used a digital camera to image 630 natural scenes, converting to CIEXYZ and then CIELUV assuming PAL CRT TV viewing conditions, primaries, and gamma function (i.e. the camera was not calibrated). A sample of pixels plotted on CIELUV axes shows that pixels are not spread out uniformly in CIELUV: there is restricted coverage of the space, and most pixels are clustered around the origin (figure 3.4). However, the images were uncalibrated, which means the accuracy of the estimates of the colors is likely to be low. Also, we do not know any details about the scenes.

The results described so far indicate that most natural scenes generally contain only a few colors, and that the bulk of the color lies in a small range compared to PCS. However, we want to know in detail the distribution of colors in a large number of natural images, in a form relevant to our modeling work. By performing a similar analysis to Yendrikhovskij (2001a), but including a histogram of pixel counts, the study of Belpaeme and Bleys (2009) comes closer to our requirements (figure 3.5). Unfortunately, this study did not use calibrated images, so the accuracy of color estimation is unclear. Also, the histogram analyis in CIELAB views the distribution of pixels on the a* and b* axes separately, but we would like to see the distribution on these axes together to learn about perceptual color. However, we can see that for natural scenes, the histograms of CIELAB a* and b* are both strongly peaked at 0, so the colors are generally of low saturation and occupy a small part of PCS.

In summary, we have seen that natural scenes appear to have a distribution of colors covering a limited range of our perceptual space, and the distribution is not uniform. There are lots of yellows, greens, and blues (from earth, vegetation, and sky). Some colors appear very rarely, if at all. Also, we have seen that the colors depend on the scene type (i.e. the reflectances present) and also the illuminant (for scene radiances). The illuminant typically gives a bias along a bluish-yellowish axis (although it can be more complex, e.g. because of indirect light reflected from or filtered through objects such as leaves). Different scene types have different reflectance spectra biases: arid desert appears yellow because of its biased reflectance spectrum, while lush foliage appears green because of a different bias in its reflectance spectrum. However, these scenes too, overall, typically cause color to vary along a bluish-yellowish axis (even without the sky; Webster and Mollon 1997).

So far, although we have seen there is clearly a bias in the distribution of colors in natural scenes (in LMS cone space, in cone-opponent space, and in perceptual space), how strong is the bias over an individual's lifetime, and what effect might such a bias have on developmental models of the type introduced in the previous chapter? To
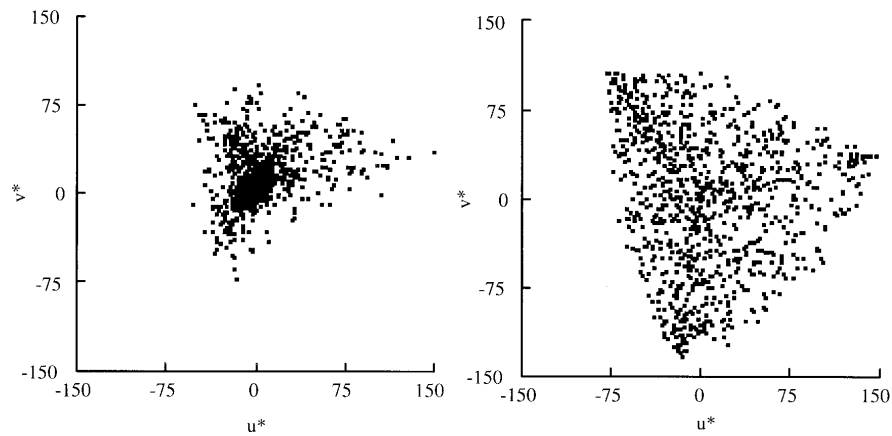
Figure 3.4: *Color in natural scenes: Yendrikhovskij (2001a). Left: an unknown sample from 630 uncalibrated natural images of unknown size plotted in CIELUV. In agreement with previous work, we see that chromaticities are concentrated in a relatively small region, although there is a larger range of chromaticities than previously reported. However, the scene types and conversion accuracy in this study are not clear. Right: For comparison, 100 random RGB values were converted to CIELUV (PAL TV parameters were used in conversion), emphasizing that the gamut of the natural scenes was restricted compared to the colors a TV can produce. Images reprinted from Yendrikhovskij (2001a).*

answer these questions, we need to find a large collection of recordings of natural scenes, containing both spatial and color (spectral) information. For such a collection, we need to determine the cone statistics as well as the distribution of hue percepts elicited, in order to understand what the effect might be on a developmental model that is learning from this input. Already, we suspect that the model will not receive a uniform sampling of perceptual color space, or uniform cone statistics. However, we need to quantify these non-uniformities.

3.3 Background 2: Databases of natural images

As discussed earlier, input to the eye is an SPD at each photoreceptor, but capturing and using this as input to a model to simulate an individual's visual experience is impractical. We need both spatial and spectral information (so therefore cannot use radiances captured by e.g. a spectroradiometer), and we need a high enough number of samples of natural scenes to be representative of visual experience. One option is hyperspectral imaging, which relatively finely samples the SPD at each pixel. A second option is to use an imaging device that has at least three sensitivities, the output from which can subsequently be converted to human cone excitations. An example
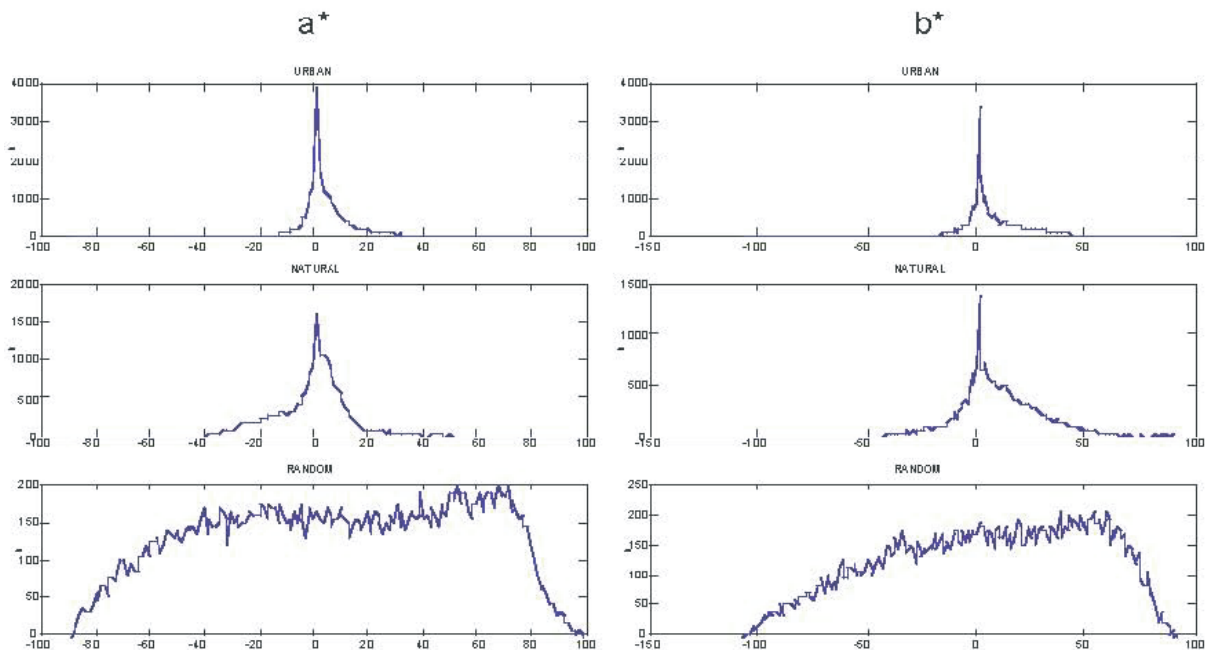
Figure 3.5: *Color in natural scenes: Belpaeme and Bleys (2009). Distribution of CIELAB a* and b* values. Top row: 25,000 randomly sampled pixels from 300 uncalibrated digital camera images of indoor and outdoor urban scenes (buildings, people, urban activities). Second row: 25,000 randomly sampled pixels from 300 uncalibrated natural images gathered from the Web (animals, flowers, and landscapes). Third row: 25,000 random RGB triples. Both the urban and natural scenes have many unsaturated pixels, and the range of CIELAB values is restricted. This agrees with previous studies. The urban images have a narrower distribution of colors than the natural scenes, but we do not know the image contents, and the conversion accuracy in this study is not clear because the images were uncalibrated. For comparison, the third row shows that random RGB images converted using TV parameters have a much wider distribution of colors. Images reprinted from Belpaeme and Bleys (2009).*

of such a device is a calibrated RGB camera: one for which the spectral sensitivities and any response nonlinearities are known. There are databases of both hyperspectral and calibrated trichromatic databases, which we will review below (summarized in table 3.2). Unfortunately, while many more images are available on the Web, the response properties of the cameras used to collect the images are very rarely known, so it is not possible to estimate cone excitations accurately enough.

Since we will be using the images as input to models in subsequent chapters, it is important that the images are freely available without restriction on re-use or re-distribution (other than attribution). Otherwise, the modeling results will not be open and reproducible.

3.3.1 Hyperspectral databases

Hyperspectral imaging was outlined in section 3.2.1. While there are several hyperspectral databases of indoor images, there are not many databases of outdoor scenes because of the difficulty in controlling for movement (and illuminant changes) between images at each wavelength. The two databases that fit our criteria are detailed in the last two rows of table 3.2, and are described below:

- Bristol Hyperspectral Images Database (BHID; Párraga et al., 1998): 29 hyperspectral images of foliage and plants (plus some sky). From Bristol, UK, in the autumn/winter, around midday. Available as reflectances plus scene illuminant.

- Foster Hyperspectral Image Database (FHID; Foster et al., 2006): 9 hyperspectral images of rural scenes: rocks, trees, leaves, grass, and earth. From Minho, Portugal, in the summer, around midday (clear sky, direct sunlight). Available as reflectances plus scene illuminant.

- Cronin and Chiao Hyperspectral Image Database (CCHID; Burton and Moorhead, 1987): 12 low-resolution hyperspectral images of natural scenes, including a target of known reflectance.

3.3.2 Calibrated trichromatic databases

Calibrated trichromatic imaging was introduced in section 3.2.1. There are several databases of calibrated trichromatic natural images, differing in trichromatic device used to capture (imaging colorimeter vs. RGB camera), and method of conversion to device-independent and LMS color spaces:

Database	Num. images, image size	Supplied formats	Relevant categories
MCID	648 768×576	RGB, SP-LMS	Animals; Flowers; Foliage; Fruit; LandWater
BCID	250 756×1134	CIEXYZ, SP-LMS, SS-LMS	Natural Objects 1, 2, 3, 4
UTCID	1204 4284×2844	RGB, CIEXYZ	Sets 1-9
UPCID	≈ 5000 1519×1007	SP-LMS	
DDB	52 1000×1280	nonlinear SP-LMS	All
ICID	25 1000×1280	CIEXYZ	Plants, flowers
FHID	9 $\approx 640 \times 820$– 1340×1020	$33 \times 10\,\text{nm}$ bands	2002 scenes 1-4; 2004 scenes 1-5
BHID	29 256×256	$31 \times 10\,\text{nm}$ bands	All
CCHID	12 128×128	$43 \times 10\,\text{nm}$ bands	All

Table 3.2: *Databases of calibrated color natural images. These databases all contain spatial and spectral information in a form that allows human cone excitations to be estimated. Hyperspectral images (last three rows) consisting of multiple bands allow most accurate and flexible reconstruction, but unfortunately are not available for a wide range of scenes. The other databases were all captured using calibrated trichromatic devices, typically RGB cameras. Depending on several factors, the accuracy of conversion to cone excitations is variable for such databases, but is accurate enough for our purposes. All databases are of scenes selected by researchers, deemed to fit into a particular category, rather than being collected by e.g. recording the experience of an individual animal. Database acronyms are given in the text.*

- Barcelona Calibrated Image Database (BCID; Párraga et al., 2010): 250 images of natural objects (foliage, fruit, flowers, sky) from around Barcelona, Spain, in the summer daytime. Sample thumbnail images are shown in figure 3.6. Conversion to LMS (or CIEXYZ) was by minimizing the error in a polynomial mapping of camera responses to LMS (or CIEXYZ) values calculated from a selection of known reflectances. Conversion accuracy was verified by comparison of predicted chromaticities to ones measured by spectroradiometer. A grey sphere of known reflectance and chromaticity is present in almost every scene, allowing the illuminant to be inferred.

- McGill Calibrated Image Database (MCID; Olmos and Kingdom, 2004): 648 images of various natural scenes and objects (animals, flowers, foliage, fruit, and land/water), including different times of day. Sample images are shown in figure 3.7. Conversion to LMS was using a linear (3×3 matrix) transformation of camera spectral sensitivities into LMS cone spectral sensitivities, with certain constraints from psychophysics ($L(580\,\text{nm}) \approx M(580\,\text{nm})$ for maximum sensitivity of the L-M cone-opponent process to be at 580 nm, and $S(506\,\text{nm}) \approx L(506\,\text{nm}) + M(506\,\text{nm})$ for maximum sensitivity of the $S - \frac{(L+M)}{2}$ cone-opponent process to be at 580 nm). The conversion is likely to be less accurate than for the BCID, because it is based only on a linear conversion of the camera spectral sensitivities to monochromatic light. Additionally, there is no validation of the accuracy for known samples, and no object of known reflectance/chromaticity in the images.

- UPenn Calibrated Image Database (UPCID; Tkačik et al., 2011): Around 5000 images from tropical savanna (Botswana). Conversion to LMS was similar to the McGill Calibrated Image Database; Olmos and Kingdom, 2004 (MCID), using a linear transformation of camera spectral sensitivities into SS-LMS cone spectral sensitivities, with the linear transform found by least squares (error minimization between SS-LMS sensitivities and the reconstructed-from-camera-RGB SS-LMS sensitivities). Additionally, LMS isomerization (excitation) rates were estimated (i.e. using standard properties of the eye). Accuracy of the conversion was checked for 24 patches of a Macbeth color checker (known chromaticities).

- UTexas Calibrated Image Database (UTCID; Geisler and Perry, 2011): 1204 images of natural scenes (earth, trees, sky, foliage, grass, etc.) from around Austin, Texas, in the summer daytime. Camera performance data are available with the database so data can be converted to both LMS and XYZ.

- IPL Calibrated Image Database; Laparra et al., 2012 (ICID): 25 close-up, indoor

images of plants and flowers illuminated by CIE D65 (northern European day-
light) and separately by CIE A (tungsten bulb light). Images were taken with
a calibrated image colorimeter (where three sensor types are calibrated to give
XYZ values).

- Doi DB (DDB; Doi, 2003): 52 images of flowers, fruit, foliage, sky, and earth.
  Converted to LMS from RGB by linear transform, minimizing errors for cone
  responses to 170 natural scene reflectances. Provided only as nonlinear cone re-
  sponses, the LMS cone excitations having been converted to nonlinear responses
  via a squashing function.

If more hyperspectral images were available, they would be the preferred source for
our modeling work. Hyperspectral images allow high accuracy, but also high flexibil-
ity (e.g. it would be simple to simulate the photoreceptor absorptions of other species).
However, very few hyperspectral images are available. Another consideration is that
in the future we wish to model more features than only OR and CR, for instance
adding motion (see chapter 6). In the foreseeable future, calibrated trichromatic video
seems far more likely than hyperspectral video. Among the trichromatic databases, the
Barcelona Calibrated Image Database; Párraga et al., 2010 (BCID) database provides
the best balance of number of scenes with calibration and conversion accuracy. The
BCID images are generally of "uninteresting" scenes, however, and so for comparison
we will also analyze the MCID, which has more images, and many images of "inter-
esting", colorful objects—of the remaining calibrated color databases, it appears to be
the most different from BCID by visual inspection.

While previous analyses have considered cone statistics for collections of images, and
have considered colors present in a variety of individual images, there is no previous
suitable bulk analysis of the colors present in large collections of images in the form
required to understand how the distribution will affect the kind of models introduced
in the previous chapter. And specifically, there has been no such analysis for the cali-
brated color image sets we could use for modeling, such as MCID and BCID.

3.4 Methods

For the MCID and BCID, we wish to find basic cone excitation statistics: total exci-
tation produced in each cone type, the correlation between cone types, and how the
excitations are distributed. Additionally, we want to know about the hue percepts gen-
erated by these natural images, using a simple (and relatively naive) model of adult

Figure 3.6: *Sample images from BCID. Random sample of 180 out of 250.*
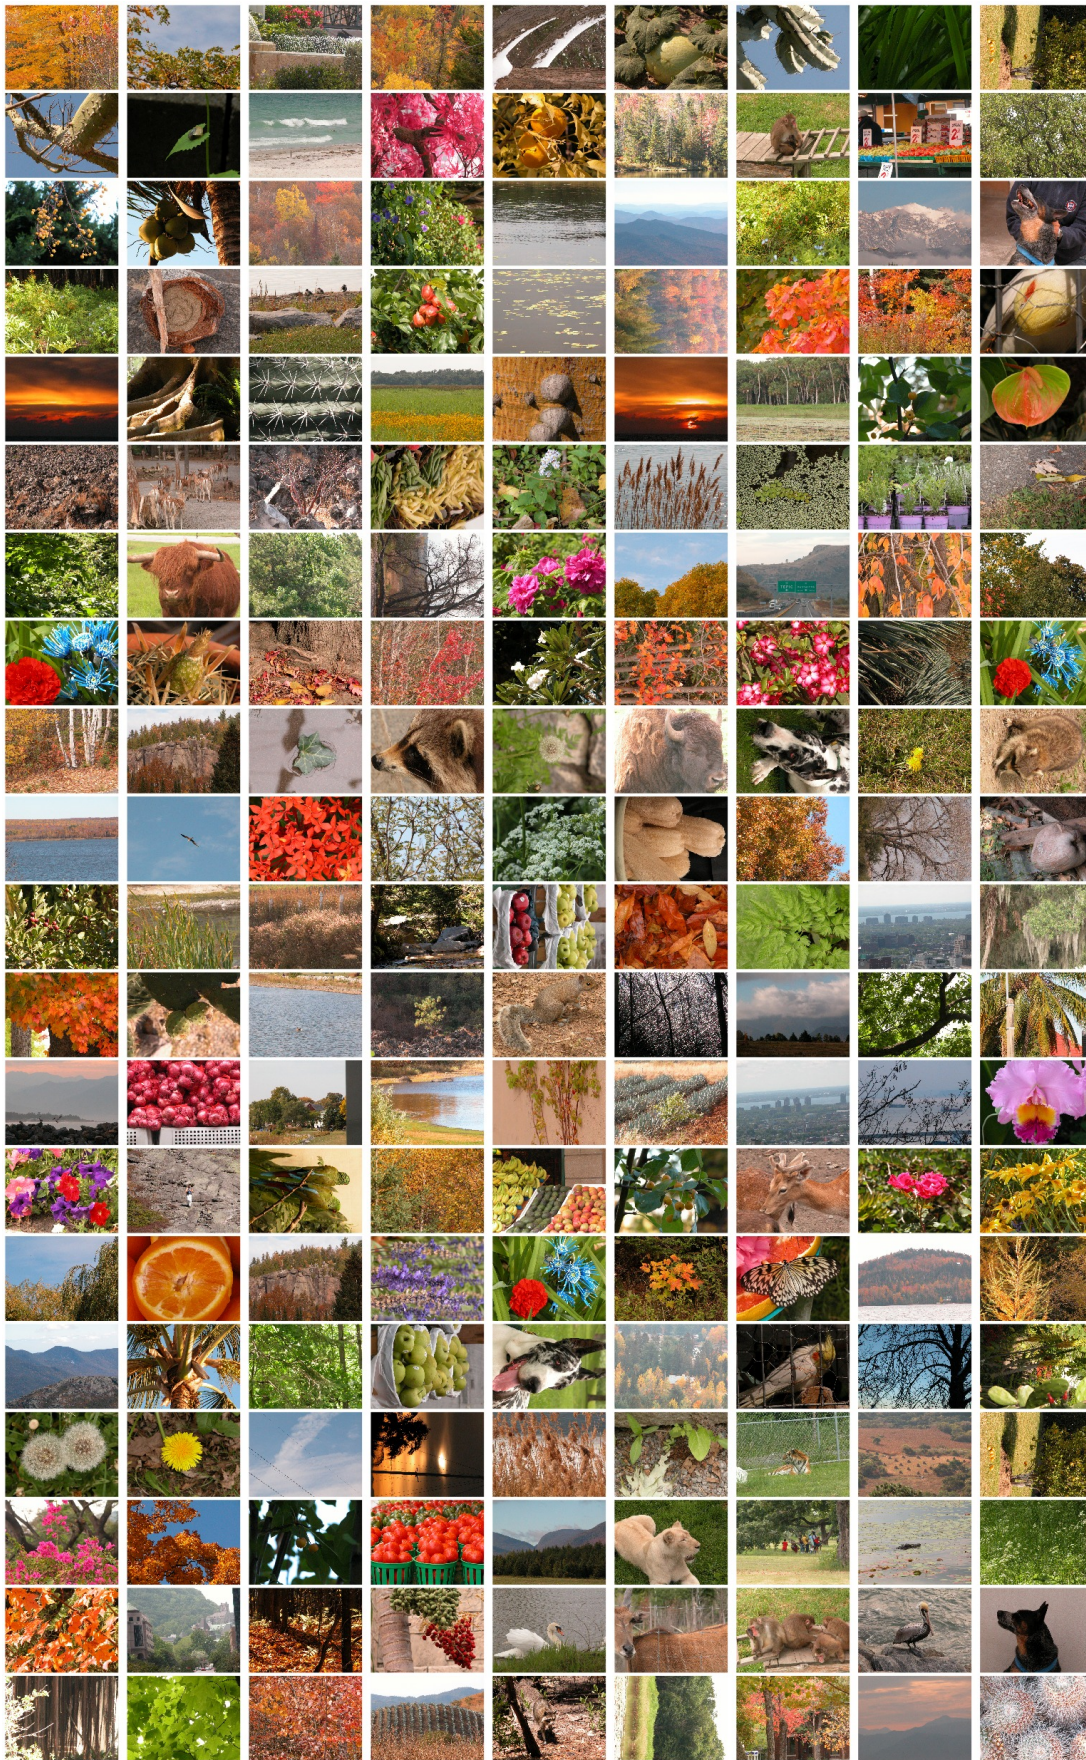
Figure 3.7: *Sample images from MCID. Random sample of 180 out of 648.*

color appearance. We want to know these things because our model will be learning from these image collections, and it will be affected by each of these aspects of the images. Therefore, we first require estimates of human/macaque cone excitations, and we require the images in the device-independent color space CIEXYZ, from which CIELAB may be computed. Once we have data in those two spaces, we can then measure the required statistics and distributions.

3.4.1 Cone excitations

We will use the Smith-Pokorny LMS $2°$ cone fundamentals (SP-LMS) as our cone space. The BCID is supplied in this format, and the MCID is supplied with Matlab code to convert the RGB images into this format. The conversions from camera images to cone excitations were outlined earlier (and are described in detail with the databases), and include an assumption about how to normalize the excitations (e.g. equal response in each cone for equal-energy white light) that is fixed across the whole database. We will discuss in chapter 4 how such fixed normalization is unlikely to be biologically realistic, and additionally does not take account of e.g. highly variable cone ratios.

With the (L, M, S) cone excitations available at each pixel of each image, we compute the following for 50 randomly selected images:

- Mean and standard deviation of the L, M, and S values over all pixels.

- Correlation coefficient of (L,M) values across all the pixels; repeat for (L,S), and (M,S) also.

- Histogram of L values; repeat also for M and S values.

- Two-dimensional histogram of (L, M) pairs; repeat for (L, S), and (M, S) also.

- Cone-opponent values $L - M$ and $S - \frac{L+M}{2}$, and a two-dimensional histogram for these also.

All computations were performed using standard functions `mean()`, `std()`, `corrcoef()`, `histogram()`, and `histogram2d()` from NumPy version 1.4 (Numerical Python; Dubois et al. 1996; Oliphant 2007).

3.4.2   Perceptual color space

We use CIELCh$_{ab}$, the polar form of CIELAB, as the perceptual color space. CIELAB is a non-linear transformation of (device-independent) CIEXYZ, and involves the creation of opponency (of CIEXYZ values, not of LMS), then a cube root (standard formulae in Wyszecki and Stiles, 1982). Additionally, in section 3.5.4 we will also use HSV color space (HSV; Smith, 1978), which is a non-linear transformation of RGB (Smith, 1978). We use sRGB, a standardized linear transformation of CIEXYZ (Stokes et al., 1996). Therefore, the main requirement is to obtain the images in CIEXYZ. The BCID is supplied in CIEXYZ format (conversion outlined in the background section, and described in detail with the database). For the MCID, we obtained tabulated camera spectral sensitivities (Fred Kingdom, personal communication) and transformed to CIEXYZ using the same method used to convert the database to LMS (described earlier in section 3.3.2), but using CIEXYZ color-matching functions (cmf) rather than LMS cone fundamentals.

Along with CIEXYZ values, to convert to CIELCh$_{ab}$ we need the whitepoint, which we take to be CIE D65, i.e. assuming the illuminant spectrum is that of northern European daylight. However, in a later section of the results, we will check the accuracy of this assumption by estimating the illuminant for a subset of the BCID images, since they include a grey ball of known chromaticity. To extract the whitepoint in each image, we located the region of the grey sphere reflecting the most light.

Once in CIELCh$_{ab}$, across all pixels of all images in the database, we record a histogram of h$_{ab}$ weighted by C$_{ab}$. h$_{ab}$ varies circularly from 0 to 1, so we will present the histogram on polar axes.

3.4.3   Hue jitter

In a later results section, we will be performing an operation we refer to as "jittering the hue distribution" of an image database. To jitter the hue distribution of a set of images, we rotate every image's hue eight times by varying amounts to create eight rotated-hue versions of the original image. Rotating an image's hue simultaneously changes every hue present in an image by the same amount. The amount of rotation for each image is drawn from a uniform random distribution, the upper bound of which depends on the amount of jitter J specified for jittering the database: $h'_{ab} = h_{ab} + U(0, J)$. For $J = 0$, the images will always be unchanged. For $J = 1$, the hue of an image can be rotated to any other value (since h$_{ab}$ varies circularly between 0 and 1). Across a collection of

images starting with a narrow h_{ab} distribution, this procedure allows the distribution to be broadened. Sample hue-jittered images are shown later in figures 3.16, 3.17, and 3.18.

Where this procedure is performed, it is done for a subset of the BCID images, Natural01 (71 images). h_{ab} is computed for each pixel using a whitepoint estimated from the grey sphere (as described above). The amount of jitter is determined by NumPy's random.uniform$(0, J)$.

3.5 Results

In the following chapter, we will be modeling the early visual system, and in particular the development of CR and OR selectivity in V1. We have chosen practical sources of input for the model, and in this section we will analyze the cone statistics, cone-opponent statistics, and color percepts generated by these input sources. This will allow us to understand what the input's effects on the model may be. Previous sections have reviewed existing work showing that color representation at all these levels is biased; in this section we expect to confirm and quantify this for the MCID and BCID.

Previous work (section 3.2) has shown that a single scene tends to have only a few colors, with different images having quite different distributions of color depending on the contents (e.g. if sky is present, or if fruit is present). To illustrate this, and introduce our analysis format, figures 3.8 and 3.9 shows sample images along with the corresponding distributions of cone and cone-opponent activations, and perceptual color histograms. Cone excitations are correlated for each image, particularly L/M, and on cone-opponent axes are distributed in a limited space. An important point to note here is that we are not saying these are the biological cone responses we would expect to measure in an animal observing this scene (if it were possible to measure such responses in a live system). The photoreceptors surely adapt, both instantaneously (e.g. von Kries gain control) and based on previous visual history. However, the only model reviewed in the previous chapter to actually use realistic LMS cone sensitivities (De Paula, 2007) did so "naively". That is, it used LMS cone responses converted from camera images, without any form of either per-image or long-term adaptation. The second stage, cone opponency, also naively used these cone responses without any adaptation. We too will begin modeling from this point, so we need to understand the statistics from this point of view.

In terms of perceptual color, we see that each of the images only contains a few colors.

Database	L : M : S	Std L, M, S	Corr L/M, L/S, M/S
BCID	1.0 : 0.82 : 0.36	0.13, 0.11, 0.077	0.99, 0.77, 0.82
MCID	0.98 : 1.0 : 0.78	0.20, 0.21, 0.20	0.99, 0.91, 0.94

Table 3.3: *Basic LMS cone statistics for natural image databases. First column: Overall cone excitations are not equal between classes. The databases make different assumptions in normalizing cone excitations, although typically it is at least assumed that a white (flat-spectrum) input should cause equal excitation in each cone class. Within a database, the normalization used is the same across images (i.e. no adaptation over time). Third column: Inter-channel correlations are all high, but L/M correlation is always the highest. (From a random sample of 50 images.)*

The previous work we reviewed in section 3.2.3 found most natural scenes contain only a few colors, so this is not surprising. In the next two sections, we will consider these properties for entire databases, rather than individual images.

3.5.1　Cone statistics in natural image databases

First we examine simulated LMS cone excitations for the natural image databases. The first column of table 3.3 shows that most databases have highest activity in the L channel, and all have lowest activity in the S channel. The overall activities depend both on the normalization method used for each cone class, and on the distribution of energy over wavelength in the images. Typically, cone responses are normalized to be equal for a flat (equal energy) spectrum, and the normalization is fixed for the database. The MCID additionally includes criteria from psychophysics about the maximum sensitivity of the cone-opponent mechanisms. Note that in human and macaque retina, the ratio of L:M cones varies greatly between individuals, and the S cones constitute only around 5 % of the total (around parafoveal retina). Hence the total excitation ratios are in some way misleading, but an important point is that without any adaptive normalization, the output from one cone class would dominate.

The third column of table 3.3 shows that all the channels are highly correlated, with the L and M channel correlation always being very high (around 0.99). These correlations are similar to those found in previous studies (described in section 3.2). We saw in section 3.2 that much of the L/M correlation is from the overlapping cone sensitivities, but the image content itself does increase correlation (particularly between M and S and L and S). Figure 3.10 shows these cone correlations graphically. Figure 3.11 shows the cone excitations are skewed towards low values, indicating the need for some form of gain control or adaptation.

(a) Sample image, cone excitation ratios, and correlations

L:M:S	corr L/M, L/S, M/S
1.0 : 0.83 : 0.47	0.98, 0.50, 0.65

(b) M vs. L, S vs. L, and S vs. M activation histograms

(c) S-LM vs. L-M activation histogram

(d) CIELCh_ab histogram
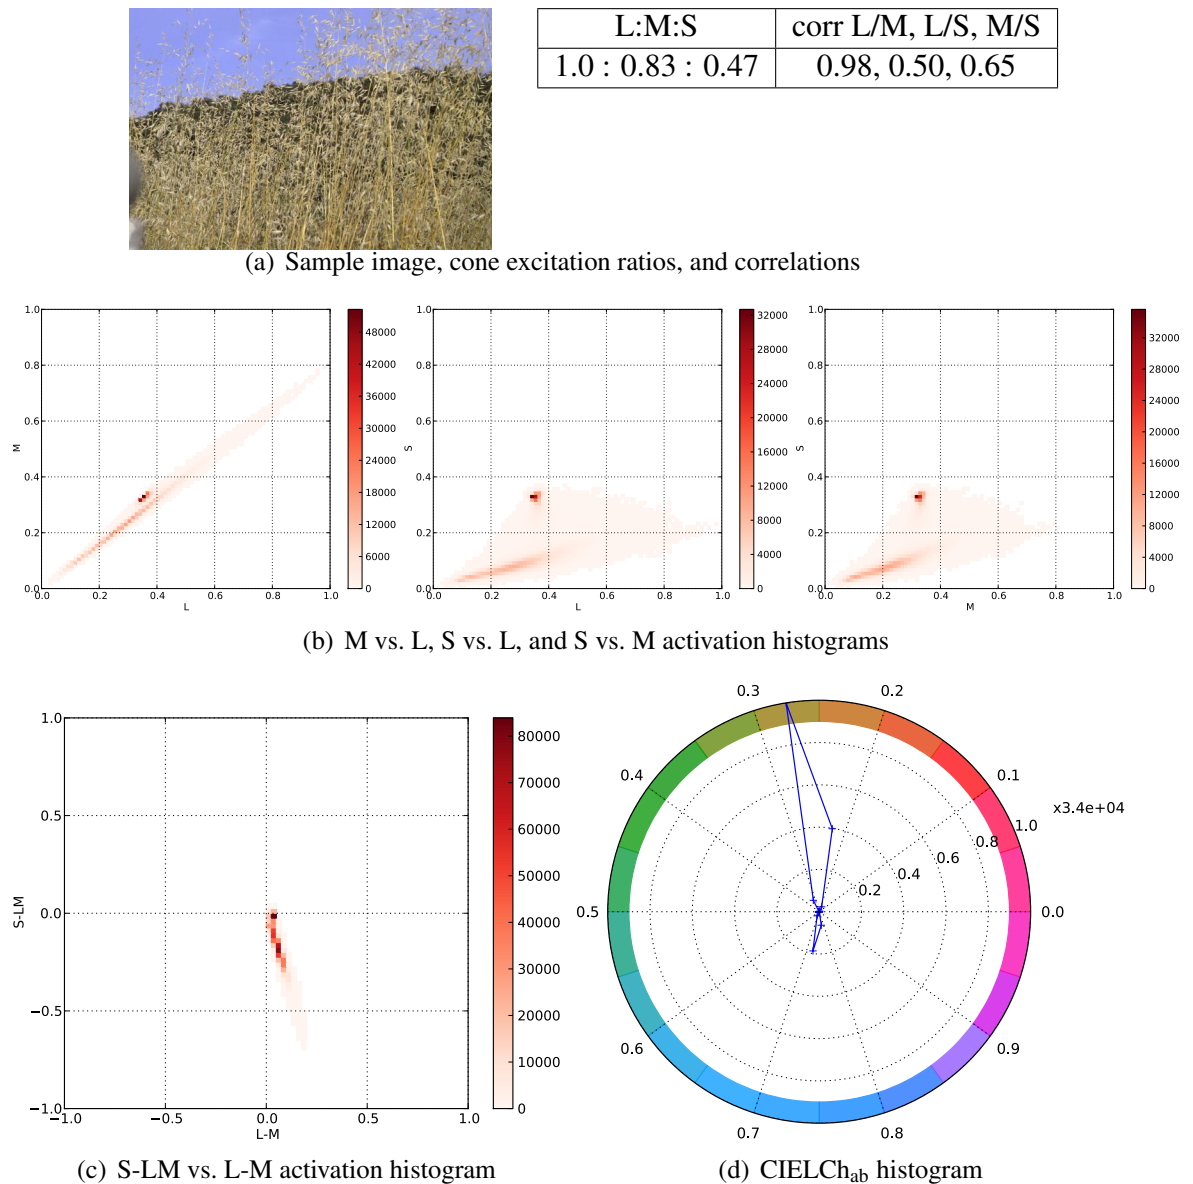
Figure 3.8: *Cone statistics and color distributions for sample image 1. (a) Image (uncalibrated RGB) and table of total L:M:S excitations and correlations over the image. The cone sensitivites are normalized assuming equal excitations to white (flat spectrum) light, and we also assume that each of the three cone types exists at each pixel. The L and M cones are highly correlated. (b) 2D histograms of pairs of cone excitations. We can again see the cone correlations, and also that most of the cone excitations occupy a small fraction of the possible area. (c) Similarly, the cone-opponent activations are biased and occupy a small area. (d) Polar histogram of CIELCh_ab hues. The polar axis is perceived hue H (from 0 back to 1), while the radial axis is a count of the number of pixels of a given hue, weighted by C. We can see the image is dominated by yellow and blue.*

L:M:S	corr L/M, L/S, M/S
1.0, 0.74, 0.27	0.97 : 0.89 : 0.91

(a) Sample image, cone excitations ratios, and correlations
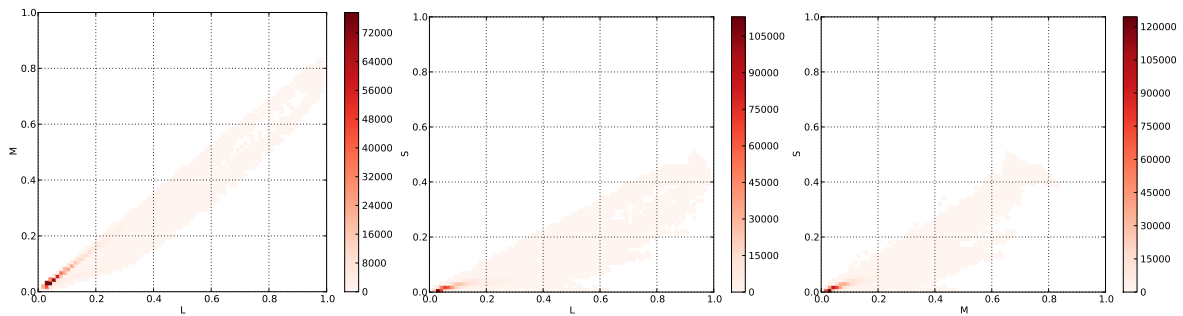
(b) M vs. L, S vs. L, and S vs. M activation histograms
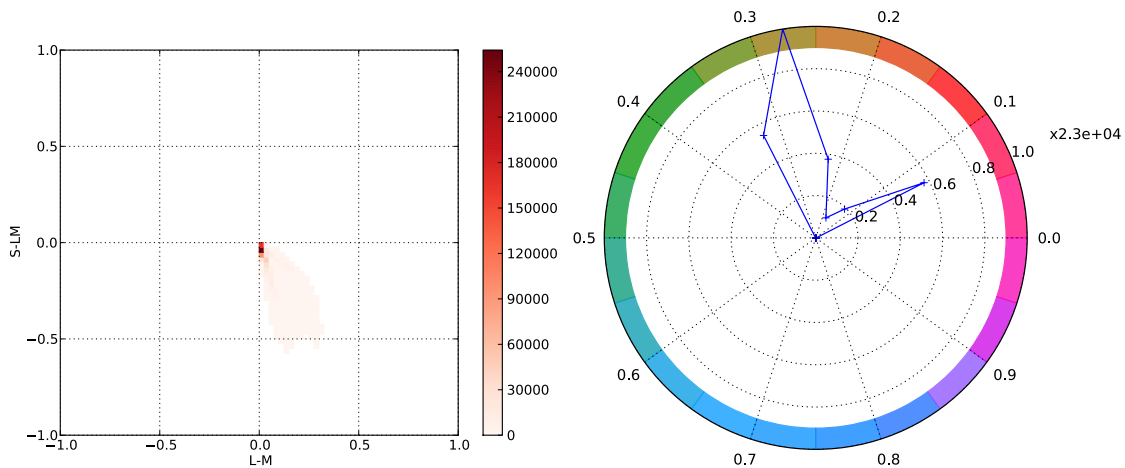
(c) S-LM vs. L-M activation histogram

(d) CIELCh$_{ab}$ histogram

Figure 3.9: *Cone statistics and color distributions for sample image 2. The plot types were introduced in figure 3.8. As with the previous sample image (figure 3.8), only a few colors (red and yellow) are present in any significant quantity, and the cone excitations are highly correlated.*

(a) BCID

(b) MCID
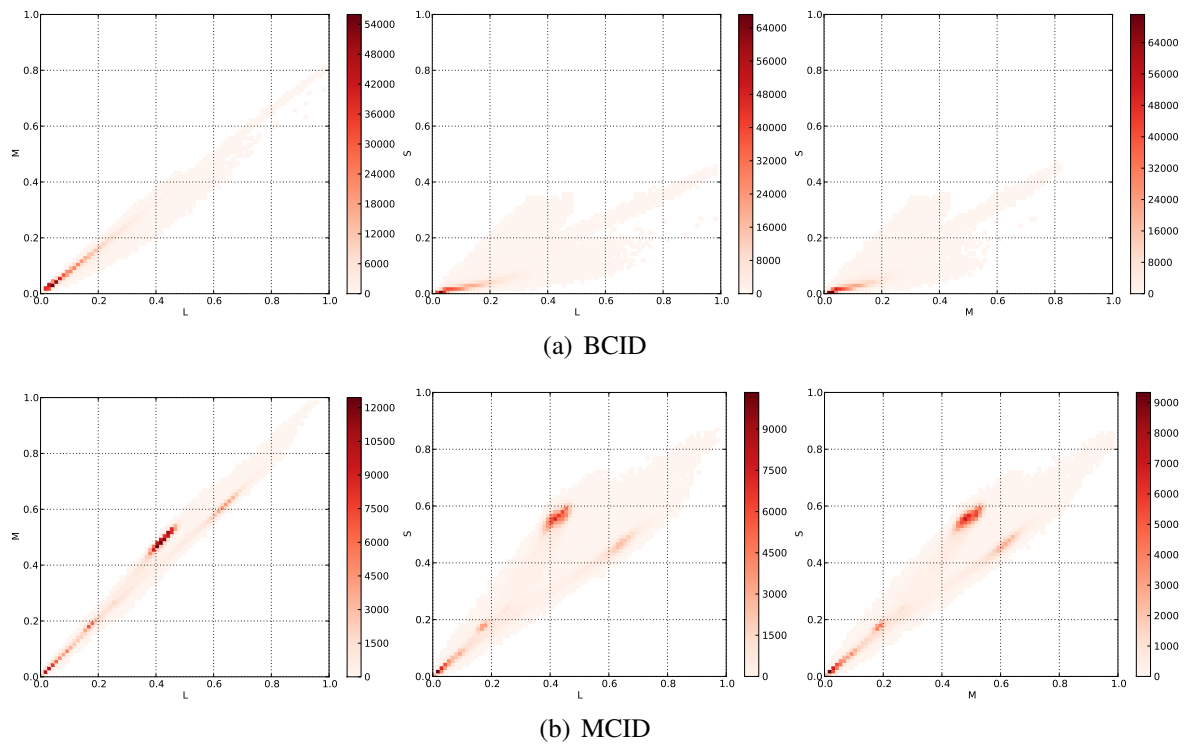
Figure 3.10: *LMS cone excitations. The format of these 2D histograms was introduced in figure 3.8. The cone excitations are highly correlated. As discussed in the text, these cone excitations do not include any variable normalization or adaptation. (From the same random sample of 50 images as in table 3.3.)*

Figure 3.12 shows the cone-opponent activations. These reflect the biases present in the cone activations: the cone-opponent activations occupy a small area, and have a bias along the negative diagonal (approximately: yellow/blue) as we saw in previous work (Webster and Mollon, 1997; figure 3.3). Even though Webster and Mollon (1997) scale their cone responses (von Kries), and the cone-opponent axes, they still find this pattern of bias, indicating that this bias is something a model of the visual system would need to deal with even with some form of receptoral adaptation present. It is also interesting to note that plotting samples from uniform color space (including CIELAB) results in a bias along the same direction, though less extreme, and less restricted (McDermott and Webster, 2012).

3.5.2 Perceptual colors in natural image databases

In section 3.2.3, we saw that natural scenes seem to contain a limited range of hues, compared to PCS, and the hues are biased along a yellow/blue axis. Here we investigate the distribution of perceptual hues for all natural images in the Barcelona and
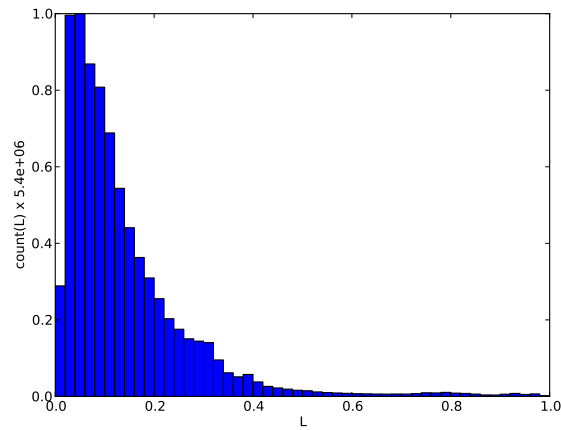
(a) L cone excitations
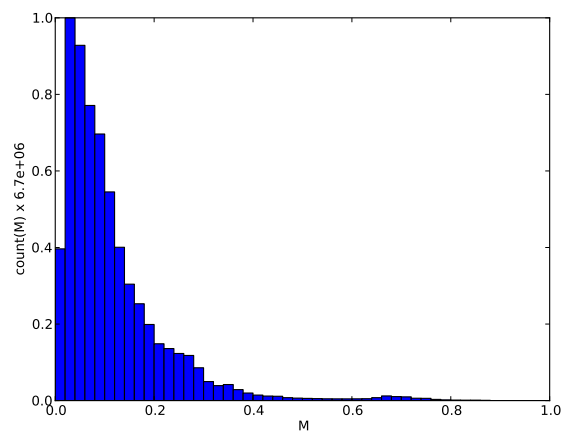
(b) M cone excitations
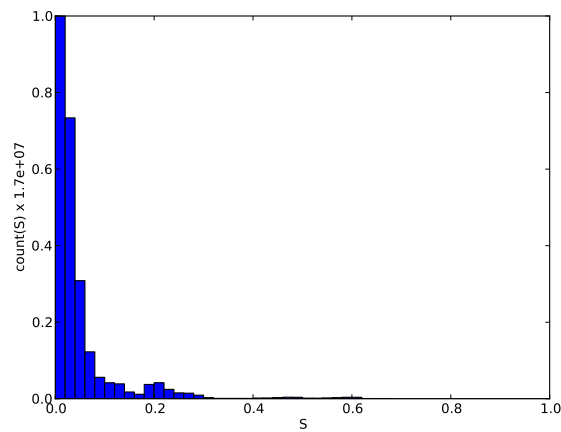
(c) S cone excitations

Figure 3.11: *LMS cone excitations are skewed. Histograms of L, M, and S cone excitations, showing more clearly the skew towards low values from figure 3.10. An unskewed distribution might be normally distributed. (From the same random sample of 50 BCID images as table 3.3.)*

(a) BCID   (b) MCID
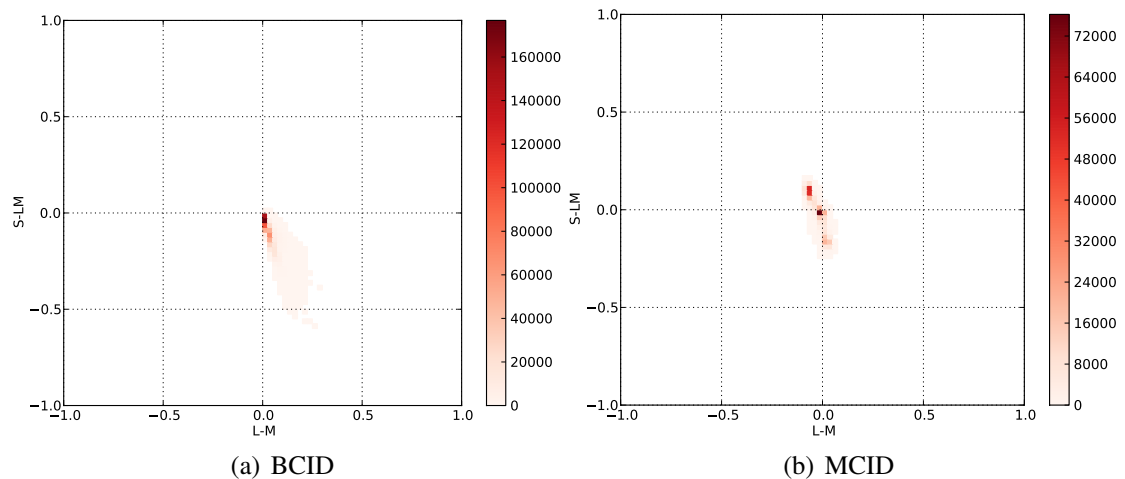
Figure 3.12: *Distribution of cone-opponent activities. Both databases show a bias along the negative diagonal -(S-LM)/(L-M), and a restricted distribution. This agrees with previous work (figure 3.3, although note unlike that figure, here we have not scaled the axes). Our model of the development of V1 will include an algorithm that learns based on the output of cone-opponent retinal ganglion cells. These results indicate that this output will be highly biased; we would expect this bias to be reflected in the model's V1. (From the same random samples of 50 images as table 3.3.)*

McGill databases. To do this, we compute $CIELCh_{ab}$ at every pixel, and create a histogram of the h_{ab} counts (weighted by C_{ab}). In subsequent chapters, we will investigate whether a Hebbian-learning–type model is able to learn PCS from this natural image input; the histogram representation should help predict the outcome.

Figure 3.13 shows that the databases contain a strong yellow/green bias, with blue also present. This fits with previous work, and reflects the dominance of foliage, earth, and sky in the scenes. Some colors, around h_{ab} 0.9, are almost entirely absent from the databases. From visual inspection, the MCID appears to be a database of interesting, eye-catching objects, whereas the BCID appears generally less visually interesting. This is reflected in the wider distribution of colors in the MCID. However, even particularly colorful subsets of MCID images, such as the MCID fruit and flowers categories, are still strongly biased to yellow/green (figure 3.13(c)).

Having characterized two databases of images that we could use as input for our modeling work, we now consider four additional questions. The first concerns the accuracy of color estimation: is our assumption of D65 illuminant reasonable? The second is whether one of the more common, simpler-to-compute color spaces in use (e.g. HSV) would be good enough to use in analyses of the model in subsequent chapters. The third and fourth questions are about the colors and statistics of the natural scenes: we

(a) BCID

(b) MCID: all natural images

(c) MCID: fruit & flowers only
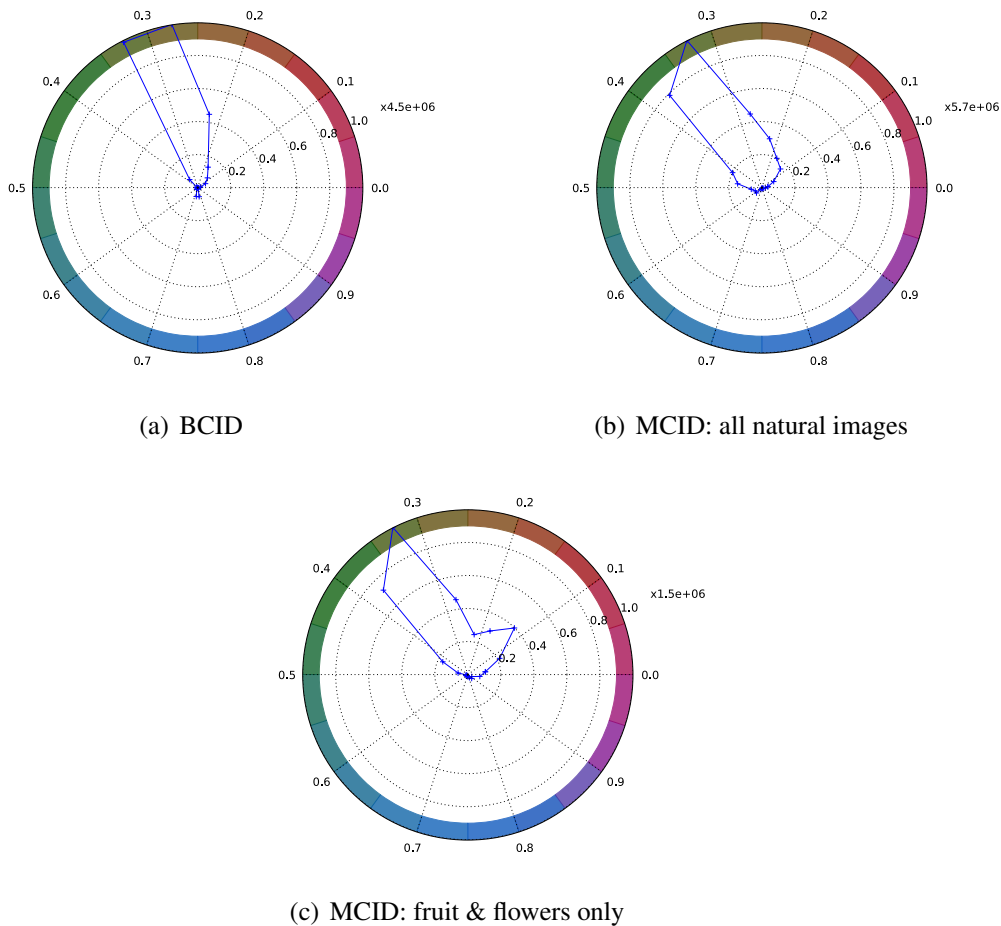
Figure 3.13: *Distribution of colors in natural image databases. CIELCh$_{ab}$ hue histograms over various sets of natural images. The histogram structure is described in figure 3.8, but shows counts of pixels binned by h_{ab} and weighted by C_{ab}. (a) and (b): For large numbers of natural scenes, the bulk of hues are around yellow/green. (c) Even using a very colorful subset of the MCID with many red objects still shows strong bias to yellow/green.*

attempt to separate the cone correlations from the distribution of perceptual colors in the images.

3.5.3   Using recorded illuminant

Light reaching the eye from a surface is a combination of the illuminant and the surface's reflectance. As discussed in the introduction, the goal of the color vision system is presumably closer to getting information about reflectance than information about the illuminant, because we usually want to know about objects, not light sources. The CIELAB color model includes an approximate normalization for the illuminant (the
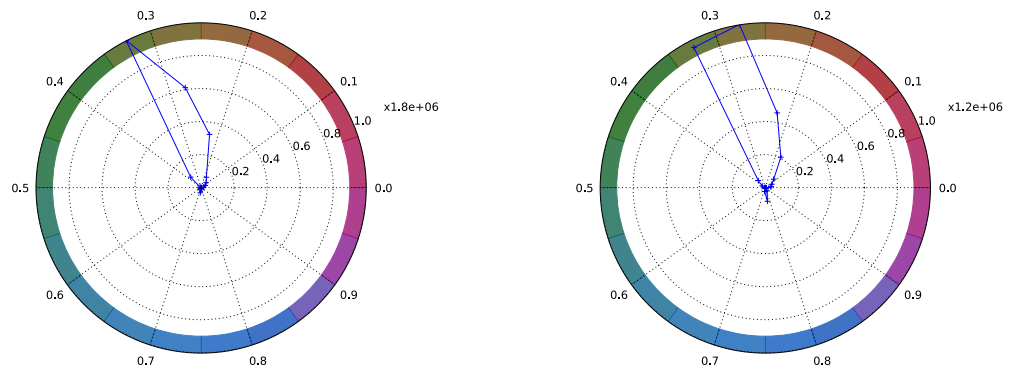
(a) Natural01 subset of BCID, assuming D65

(b) Natural01 subset of BCID, using per-image whitepoint

Figure 3.14: *Perceptual colors with more accurate whitepoint. (a) A subset of the BCID, analyzed as in figure 3.13 (assuming a fixed D65 whitepoint). (b) The same subset of images as (a), except estimating the illuminant from the lightest region of the grey sphere. This does result in a broadening of the distribution, but for these natural scenes the effect is small. In a database with images taken in many different locations and situations (e.g. under a canopy of green leaves vs. in open space), at different times of day, in different weather conditions, using an accurate per-image whitepoint could make more difference.*

"whitepoint"), intended to approximate the early visual system's short-term adaptation. When we analyzed perceptual colors in the image databases in the previous section, we assumed the illuminant was D65 (i.e. northern European daylight). This seems like a reasonable assumption, because many of the images were indeed collected under those conditions. However, the BCID includes a grey sphere (of known reflectance and chromaticity), which can be used to estimate the illuminant. Therefore, we can see the effect of assuming a fixed illuminant rather than using the illuminant from the image itself. Figure 3.14, using a subset of the BCID images, shows that using the per-image estimated illuminant does result in a slightly broader distribution of colors—but the bias is still strongly present.

3.5.4 Simpler, non-uniform perceptual color space: HSV

HSV (like the closely related HSL space) is a non-linear transformation of RGB, and is device dependent in the same way as the underlying RGB space from which it is calculated. HSV has hue, saturation, and value, but these are not independent perceptual dimensions (e.g. changing hue affects also the lightness/brightness we perceive),

(a) HSV   (b) CIELCh$_{ab}$
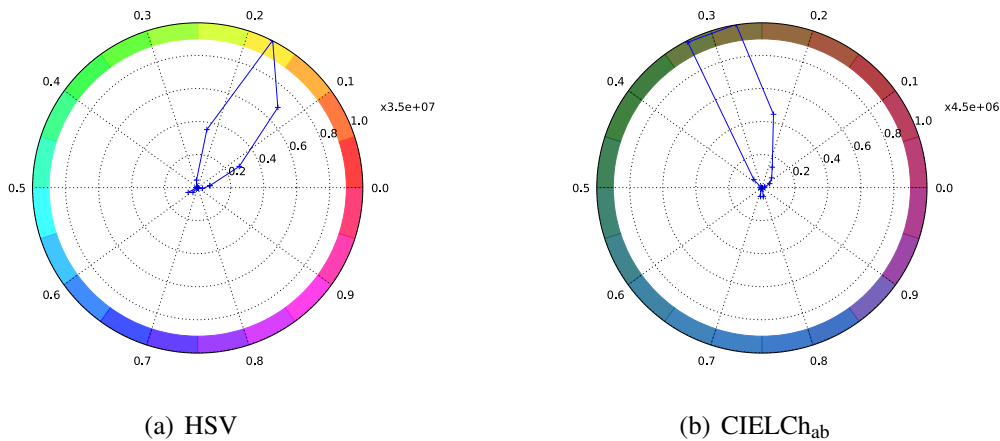
Figure 3.15: *Analyzing databases in HSV rather than CIELCh$_{ab}$ color space. Analyzing the BCID in HSV leads to the same qualitative conclusion as the CIELCh$_{ab}$ analysis, i.e. the bias is of similar magnitude, and also in a yellow direction. HSV would therefore be an acceptable space to use in our later modeling work with large sets of natural scenes.*

and additionally are not perceptually uniform (e.g. a given change in HSV hue in one part of the color circle changes the perceived color far more than in another parts). On the other hand, HSV is an extremely commonly used space in computer graphics, captures the basics of perceptual color space (e.g. red and blue are as close as red and green, and are joined by magenta) and has even been used in investigations of perceptual color representation in animals (Li et al., 2014). Therefore, it is interesting to see the HSV equivalent of the CIELCh$_{ab}$ results. We observed that for some individual images, the differences can be significant, but figure 3.15 shows that the results are quite similar over a large number of images, demonstrating that HSV is a reasonable overall approximation for these natural scene images.

3.5.5   Natural scenes with modified hue distributions

If we were able to collect a set of natural images that perfectly reflected perceptual color space, rather than the colors in real natural scenes, what would it be like? Such a set of images may be useful because it could allow us to separate the distribution of perceptual colors from the correlated LMS cone sensitivities, which could be useful for understanding our later computational modeling.

Starting from a subset of the BCID, Natural01, we "jittered" the hues present in the set of images by varying amounts in order to create collections of images with different hue distributions. The method was described in section 3.4.3, and we created three dif-

Hue jitter	L : M : S	Std L, M, S	Corr L/M, L/S, M/S
0.0	1.0 : 0.83 : 0.38	0.17, 0.14, 0.10	0.99, 0.87, 0.91
0.2	1.0 : 0.84 : 0.40	0.14, 0.11, 0.089	0.99, 0.80, 0.85
0.6	1.0 : 0.87 : 0.55	0.14, 0.11, 0.091	0.97, 0.85, 0.87
1.0	1.0 : 0.90 : 0.70	0.14, 0.12, 0.11	0.97, 0.77, 0.83

Table 3.4: *Basic LMS statistics for hue-jittered BCID Natural01 images. Jittering the hue does have some effect on the LMS correlations, though they remain high. From a sample of 50 images.*

ferent collections: jitter 0.2, jitter 0.6, and jitter 1.0. Samples of these three collections are shown in figures 3.16, 3.17, and 3.18.

Figure 3.19 and table 3.4 show the results of this manipulation. The collections still have high LMS correlations, but they are lower than for the original database. The distribution of perceptual hues of course becomes more uniform, although because the images are so biased in perceptual hue space, it takes many more than eight whole-image hue rotations before the distribution can become fully smooth (investigated further in chapter 4).

3.5.6   Symmetric sensitivities

We saw in the previous section that even when the distribution of perceptual hues is made uniform, there are still high correlations in the input signal. This is because LMS sensitivities overlap to a high degree, particularly for L and M cones. In contrast, sRGB sensitivities overlap much less, and the overlap is more symmetric. Using the full Barcelona and McGill databases (as in section 3.5.2), but converted to sRGB rather than LMS, results in lower and more equal correlations for the same distribution of perceptual hues. The correlations are shown in table 3.5. This confirms that models introduced in the previous chapter that did not use LMS sensitivities would have been learning an unrealistic input distribution. However, it also suggests that we could use sRGB images as input to separate the effect of cone correlation from the effect of perceptual hue distribution.

3.6   Discussion

In this chapter, we have seen that the colors in natural scenes appear to be a highly non-uniform sampling of PCS. Yellows, greens, and blues (from earth, vegetation,

Figure 3.16: *BCID with hue jitter 0.2. 180 out of 560 images (containing some duplicates differing only by hue).*

Figure 3.17: *BCID with hue jitter 0.6. 180 out of 560 images (containing some duplicates differing only by hue).*

Figure 3.18: *BCID with hue jitter 1.0. 180 out of 560 images (containing some duplicates differing only by hue).*

Database	R : G : B	Std R, G, B	Corr R/G, R/B, G/B
BCID	1.0 : 0.90 : 0.53	0.12, 0.10, 0.11	0.88, 0.62, 0.82
(LMS)	1.0 : 0.82 : 0.36	0.13, 0.11, 0.077	0.99, 0.77, 0.82
MCID	0.79 : 1.0 : 0.62	0.16, 0.21, 0.18	0.78, 0.60, 0.86
(LMS)	0.98 : 1.0 : 0.78	0.20, 0.21, 0.20	0.99, 0.91, 0.94

Table 3.5: *Basic RGB statistics for natural image databases. The same images converted to sRGB rather than LMS channels have lower between-channel correlations. RGB images are therefore not suitable as realistic training input for models, but could be used to separate the effect of LMS correlations from the effect of a non-uniform sampling of perceptual color space. (Note: linear RGB is used here, with no gamma applied.)*

sky) are common, while other colors rarely appear at all. We characterized this bias in two databases of images of natural scenes that we could use as input for training a developmental model of OR and CR preferences. As well as the non-uniform sampling of PCS, we have seen that the cone excitations that would experienced by a model using these databases are highly biased. First, total excitation of cone types in databases varies. Second, cone excitations are highly correlated, and not just because of cone sensitivity overlap. Third, cone excitations are highly skewed towards low values. Any model developing a perceptual color space using natural images as input would therefore need to overcome the biased distribution of perceptual colors, uneven channel activations, skewed channels, and correlated channels if it is to develop a representation of PCS. How that could happen is the subject of subsequent chapters.

Apart from learning about the issues to be faced by a model using these databases as input, a number of other questions follow from our results so far. We consider these below.

Figure 3.19 *(following page)*: *Color distributions of hue-jittered BCID images. Top panel: jitter 0.2; middle panel: jitter 0.6; bottom panel: jitter 1.0. As the jitter is increased, the distribution of perceptual hues broadens (bottom right in each panel). The distributions of cone excitations, and cone-opponent activations, also change. However, the correlations remain high (see also table 3.4), suggesting that bias in the hue distribution will not be the only problem faced by our subsequent modeling of V1. Even if the hue bias were not present, input to a model would likely still present difficulties for learning mechanisms such as Hebbian learning, particularly the highly correlated L and M signals. We will return to this issue in the following chapters. (Hue histograms over the entire 560 images; cone statistics from a sample of 50 images.)*
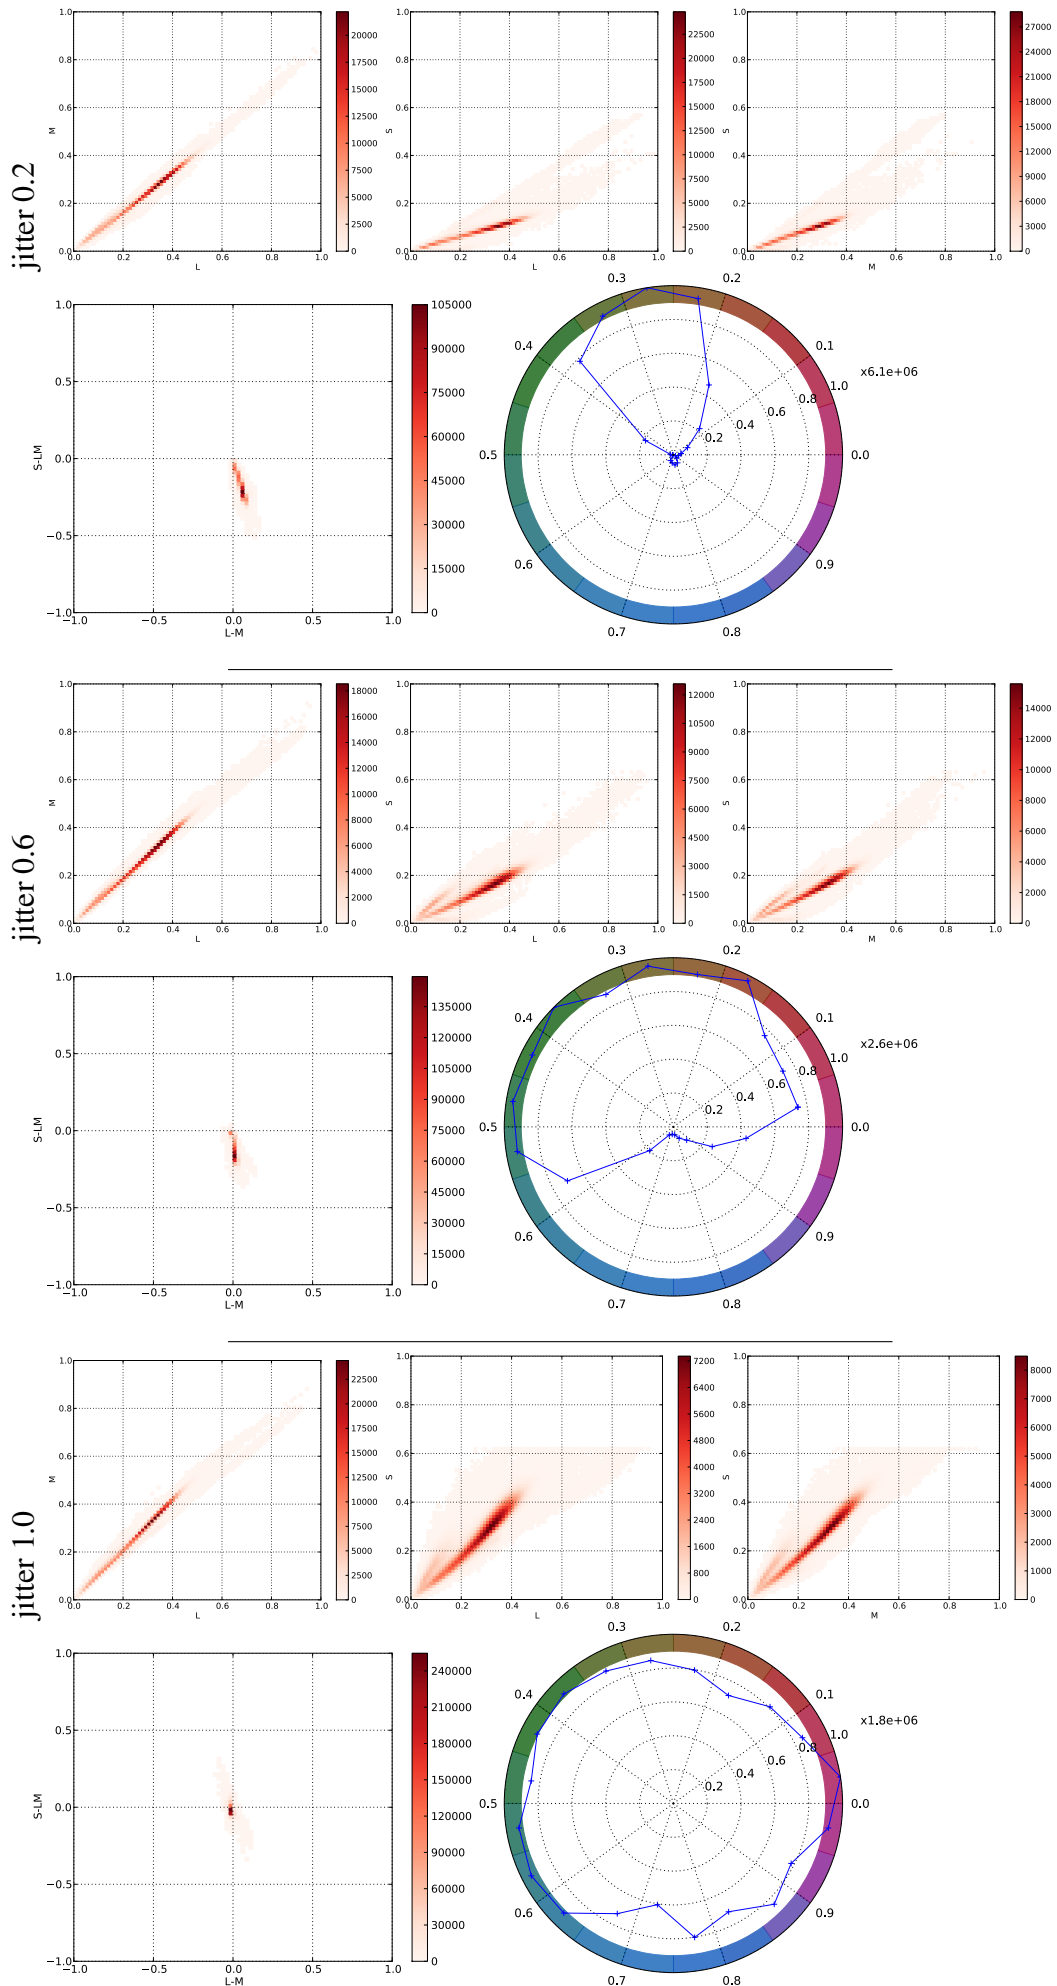
3.6.1   Do existing databases accurately reflect visual experience?

For modeling the development of PCS, it is important to have a representative sampling of visual experience. As discussed in the introduction, we decided that natural scenes probably best reflect the visual experience of macaques used in the imaging experiments, and of many humans (from whom perceptual color space is constructed). Although the color statistics of man-made scenes are probably different from natural scenes, and certainly have effects on details of PCS (e.g. the direction of minimum sensitivity predicted in CIELAB is closer to orange-cyan, rather than yellow-blue as would be predicted from natural scenes; McDermott and Webster, 2012), we believe the colors present in man-made scenes likely reflect human PCS. However, the imaged macaque monkeys could have been reared in extremely colorful or extremely drab environments, and we do not know which.

If we assume that natural scenes are the right scenes to use, the next question is have the natural image databases sampled them appropriately? The scenes are not at all a random, unbiased sample of visual experience in a way that e.g. data from a camera attached to an animal would be. The MCID has lots of interesting objects, and thus might be a good database to simulate paying attention to colorful things, while the BCID seems less deliberate. However, in both cases, we found the distribution of colors and statistics agrees with previous studies of the colors in natural scenes.

One could also argue that the databases are too small to be representative. However, it seems clear from analysing subsets of databases that collecting more images will result only an increasingly clear bias, as although more of the rare/non-existent colors may be encountered, these will be dwarfed by the increasing quantity of e.g. yellow that will also be encountered.

Finally, another factor not incorporated into the databases is that the scenes themselves (i.e. the reflectances present) likely change over time, in particular with season. For instance, Webster et al. (2007) shows there are significant differences in reflectances and cone signals in several different regions at different times of year, for a small number of scenes. While the MCID does have some images of autumn and winter, it is still largely composed of summertime images, and there are no repeated images of the same scene at different times of year.

3.6.2   Alternative viewing conditions

Apart from whether or not the natural scenes in the databases are "the right scenes", another question is about the effect of varying viewing conditions. In reality, these scenes would be experienced at different times of day and in different weather conditions. All of those things may affect the illuminant and hence the statistics of the scene radiances and therefore possibly the cone statistics and colors perceived (time of day: Wyszecki and Stiles, 1982; Lovell et al., 2005; Granzier and Valsecchi, 2014; weather: Lee and Hernández-Andrés, 2005). We mentioned earlier (and will discuss further below) that the early visual system appears to rapidly adapt to the scene illuminant, and hence this may not be an important factor for perceived hue, but this is still something a real visual system would have to deal with.

One possibility for future work would be to take a database such as the BCID containing a target of known reflectance, remove the illuminant, and then re-illuminate with simulations of the different illumination conditions (from times of day and weather conditions). Another possibility would be to capture new calibrated images of scenes at different times of day, and in different weather conditions (see the chapter 6).

3.6.3   Model of perceived color

We mentioned in section 3.2.3 that CIELAB is far from a complete description of human PCS, and not the most advanced model of color appearance. In particular, while CIELAB approximately takes account of the illuminant (via a "wrong" von Kries transform, since it is not performed in LMS space, but in XYZ; Fairchild, 2005), it does not take account of other important factors in color vision such as spatial context. Therefore, it would be interesting to repeat the $CIELCh_{ab}$ analyses using a recent, more complex, color appearance model. While not currently necessary for understanding the behavior of models of the early visual system reviewed in the previous chapter (e.g. because those models do not include adaptation), a more accurate description of the distribution of color percepts elicited by natural scenes would be important to confirm the bias found in $CIELCh_{ab}$ space, and hence for understanding the human visual system.

3.6.4 Normalization, adaptation, saliance

We have presented distributions of cone and cone-opponent activations, as well as perceptual hues. In doing so, we have not considered any neural normalization or adaptation, except in our model of perceptual hue (in which an estimate of the adapting effect of illuminant is taken into account). The color visual system is known to adapt over both short and long timescales, at multiple levels. The reason for naively considering the cone and cone-opponent statistics (i.e. not including any adaptation) is that this is what our initial model will do, since it will be based on the previous work reviewed in chapter 2. However, to develop a proper understanding of the representation of color at various levels of the visual system, it will be important to include normalization and adaptation. Not doing so may account to some extent for the bias in color statistics, although previous work indicates that strong biases are still present even after simulating normalization and adaptation (Webster and Mollon, 1997).

Another consideration from the biological point of view is higher-level mechanisms, which may cause the salience of stimuli to vary. While we may encounter far more of some colors than others, perhaps a higher-level mechanism causes more weight to be given to rarer stimuli (McDermott et al., 2010). This is also something we have not considered here, because it would greatly complicate the modeling process.

3.7 Conclusion

We have analyzed two calibrated color image databases (MCID and BCID), showing that the databases contain a distribution of colors with a strong bias towards yellowish hues, and additionally have highly correlated cone excitations, with excitations highly skewed towards low values. This suggests that any input-driven model of the development of color vision using these image databases as input will have to deal with these differences between the input and color space. Our review of previous work indicates that these biases in the color distributions in the natural scenes is not at all specific to these two databases, but a general result.

Chapter 4

Naive model of color map development

4.1   Introduction

The main goal of this thesis is to understand how wiring for perceptual color space (PCS) arises in the brain. As previous chapters have discussed, how this happens is not obvious. In chapter 2, we saw that PCS is quite different from wavelength space (WS). For instance, the most separated visible wavelengths in WS—around 400 nm and 700 nm—appear more similar to each other (red and violet) than either does to a middle wavelength—around 550 nm—which appears green. We also saw that visual experience is critical for the development of color vision, but at the same time our PCS is relatively insensitive to a wide range of environmental and physiological variations. If our visual experience were a uniform sampling of PCS, a learning system could easily construct PCS out of WS—but visual experience is a highly *non*-uniform sampling of PCS (chapter 3). Furthermore, as covered in the previous two chapters, the non-uniform sampling is likely different for different locations, times of day, seasons, and weather conditions, yet PCS remains relatively stable between individuals and for one individual over time. Additionally, PCS is reasonably consistent between different individuals despite quite different physiologies (e.g. large individual differences in L:M cone ratio; chapter 2), and for one individual despite changes in physiology and anatomy over time (e.g. the lens yellows with age). So, where does PCS come from—how is it constructed in the brain?

We can start by identifying where in the brain PCS first appears—giving us some bounds on the biology involved in the transformation, even if we do not yet know

all the processes involved. We mentioned above that input to the early visual system is certainly different from PCS; previous chapters have also reviewed how the early visual system does not appear to be organized according to PCS at the retina or LGN stages. However, there is evidence for organization by PCS in primary visual cortex (V1). These experimental results come from a small number of macaques, and studies have used various experimental methods and have made varied claims, so the results about hue representation in V1 are not yet completely established. Additionally, the experimental work done so far does not address how wiring for color (CR) could arise. This situation motivates our computational modeling.

In chapter 2, we saw that computational modeling has already provided some insight into the development of V1. We reviewed previous models that explored how wiring for orientation (OR) and CR organization could develop, and that began to explore how hue representation could develop. While a number of models have been used to investigate various aspects of OR and/or CR development in V1, in some cases making similar claims as each other (chapter 2), we found the model providing the closest starting point for understanding OR and CR organization was the De Paula (2007) model. This model has a realistic architecture comprising the principal stages of the early visual system: LMS photoreceptors, spatially and cone-opponent RGC/LGN processing, and a self-organizing model of V1. Results from the model match a number of experimental findings: the model develops a generally OR-selective map with patches of CR-selective cells, matching experimental results **ER1-ORmap** and **ER2-patches**.

However, the De Paula (2007) model—along with all previous modeling work—does not explain experimental results about the representation of hue in V1. We have not seen how a single CR patch can represent many or all hues, nor have we seen hues in a CR patch having a spatial organization mirroring that of PCS (experimental results **ER3-range** and **ER4-peaks**). The principal goal of this thesis therefore becomes to extend the De Paula (2007) model to develop a realistic representation of hue, and so to generate hypotheses about how appropriate wiring might arise. In doing so, we will come to understand the significance of certain previously neglected features of the early visual system's color pathway.

Before we can proceed to investigate the development of PCS in the De Paula (2007) model, we must first address some of the other problems with previous work that we highlighted in chapter 2. When we reviewed the state of the art in modeling the development of PCS in V1, we found there were a number of problems that applied to all previous modeling work. There were three groups of problems, the first being *PCS representation problems* (as described in chapter 2): problems **P1-range**, **P2-PCS**, and

P3-pref. However, the other two groups of problems are ones that prevent us from being able to model easily and compare with experimental data, so must therefore be addressed first:

- *Not directly comparable measurements*: Existing models have not been compared closely to experimental data about the representation of CR. Experimental work indicates that within a CR patch, peak responses to different hues are spatially separated, but there is no equivalent modeling analysis (problem **P4-peaks**). Additionally, modeling so far does not identify CR patches by comparing responses to red/green and black/white gratings, as experimental work does (problem **P5-RGBWgratings**), which will differ from analyses based on responses to all hues.

- *Modeling problems*: Previous modeling work does not address some of the important problems discussed in chapter 3. First, we saw that databases of natural scenes are a highly non-uniform sampling of PCS, and also when converted to photoreceptor excitations result in highly biased photoreceptor statistics. Previous models have used small selections of images and/or manually chosen L, M, and S channel scaling factors to ensure a CR preference map that represents all photoreceptor classes (problem **P7-LMSscaling**). Changing input images therefore requires an extensive parameter search, which is somewhat circularly guided by the modeler's judgment of the resulting CR preference map's appearance. Second, for some models we do not see the progress of CR map development over time, and for others we know CR map development does not proceed smoothly. In the case of the De Paula (2007) model, because the model of V1 used (LISSOM) simulates homeostasis via sudden jumps in parameter values, there can be unstable changes in the V1 preference map during development (problem **P8-V1StableDev**).

It is important to solve these problems first because they make investigating the model difficult; solving them is this chapter's goal. To do this, we will introduce a new model, which we will refer to as the naive model. We will evaluate this new model against criteria we reviewed in chapter 2 (experimental results **ER1-ORmap**, **ER2-patches**, **ER3-range**, and **ER4-peaks**). We will show a close match to experimental results for joint OR and CR organization (experimental results **ER1-ORmap** and **ER2-patches**) while solving problem groups *not directly comparable measurements* (**P4-peaks** and **P5-RGBWgratings**) and *modeling problems* (**P7-LMSscaling** and **P8-V1StableDev**). The naive model will therefore be the starting point for the rest of this thesis. However, we will see that the new model still has the *PCS representation problems* of previous

work. Later on, in the following chapter, we will investigate the *PCS representation problems* in detail and show how a model *can* develop a realistic PCS, and we will relate those conditions to problems the brain is overcoming in developing PCS from visual input.

4.1.1   Contributions of this chapter

- By simulating homeostatic mechanisms at the photoreceptor and V1 levels, we extend the current state of the art in modeling the development of OR and CR to allow (a) investigation of development of PCS without requiring the modeler to set parameters to anticipate experimental results, thus providing fewer degrees of freedom and avoiding overfitting, (b) use of larger and more varied sets of input data (as analyzed in chapter 3), likely to be more realistic, and (c) detailed comparison with experimental results. The resulting model, the naive model, therefore becomes the best available model of the development of OR and CR selectivity in V1 because it has better results (e.g. higher quality OR map), more detailed analysis, and fewer problems than the previous best model.

- We show that while the naive model is a good match for joint OR and CR organization, in a detailed comparison with experimental work it is a poor match for CR organization: the model cannot support a spatial or preference-based code for hue. We suggest that the model is sensitive to physical and environmental variabilities that the early visual system is not sensitive to, which will be addressed by models in later chapters.

4.2   Methods

The model presented here consists of three stages (figure 4.1), and is developed from previous modeling work (figure 4.2). The subcortical pathway (stages 1 and 2) is generally based on the De Paula (2007) model's subcortical pathway, but V1 (stage 3) is from the Adaptation, Laterally connected model (AL; Stevens et al., 2013b) rather than LISSOM as used in De Paula (2007). The reason for using AL is that it is similar to LISSOM but simpler and more robust, giving a realistic level of stability during development (discussed in a later section).

In overview, sheets of firing-rate–based point neurons represent the photoreceptors, RGC/LGN, and V1. At every iteration, an input is drawn on the photoreceptors, and

units in the RGC then compute their activities based on photoreceptor output; finally, units in V1 compute their activities based on afferent activation from the RGC combined with lateral connections within V1. Each iteration therefore represents one visual fixation. As the simulation proceeds through 10,000 training iterations, V1 self organizes through incremental Hebbian learning of its afferent and lateral connections. One iteration of the model is illustrated in figure 4.3

In the following sections, we will walk through the model's three stages in more detail, which necessarily includes a detailed description of the models of Stevens et al. (2013b) and De Paula (2007). Along the way, we will introduce additions to previous work in order to solve the *modeling problems* mentioned in the previous section (and more fully described in chapter 2). The additions involve simulating adaptive homeostasis in the photoreceptors and V1. To conclude this section, we will demonstrate the model in action.

Parameters are listed in appendix C, and the model and simulator used to generate results in this chapter are freely available from www.topographica.org.

4.2.1 Stage 1: Photoreceptors

The process of simulating photoreceptors is outlined in figure 4.4; here we will describe the details. During training (simulating an individual's lifetime of visual experience), input to the model comes from a database of calibrated color images of natural scenes. We start with images in CIE 1931 XYZ color space (CIEXYZ; Smith and Guild, 1931) from the Barcelona Calibrated Image Database (BCID; Párraga et al., 2010). At each iteration, a patch of CIEXYZ image is first converted to light-linear cone excitations (i.e. absorptions that are linearly related to the power of incident light). Each (X_i, Y_i, Z_i) image pixel i is converted to (L_i, M_i, S_i) using a 3×3 matrix $\mathbf{M_1}$:

$$\begin{pmatrix} L_i \\ M_i \\ S_i \end{pmatrix} = \mathbf{M_1} \begin{pmatrix} X_i \\ Y_i \\ Z_i \end{pmatrix} \qquad (4.2.1)$$

The matrix $\mathbf{M_1}$ is a linear transformation of the CIEXYZ cmf to human/macaque foveal LMS cone fundamentals (the Smith-Pokorny LMS $2°$ cone fundamentals (SP-LMS) fundamentals), although alternative transformations can be used. For some transformations, or if the (L_i, M_i, S_i) are estimated by an alternative method (see below), there may be negative values, in which case we threshold at 0.
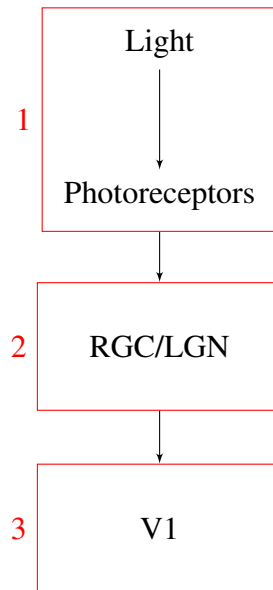
(a) Model's three stages.
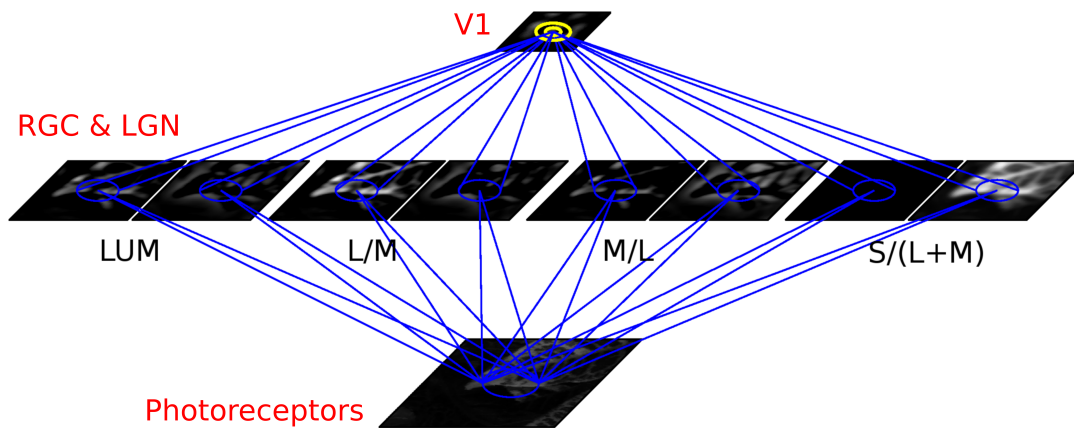
(b) Overview of model's architecture.

Figure 4.1: *Overview of naive model. (a) Schematic showing the model's three main stages. Subsequent schematic figures in the same format give more detail about each of the three stages. (b) First, the response of human/macaque LMS cones to light is simulated by processing patches of calibrated color images according to human cone spectral sensitivities. Second, spatial and cone opponent processing is simulated by the RGC, which comprises pairs of ON and OFF sheets, using Difference-of-Gaussians receptive fields. The LGN stage is folded into the RGC, assuming that they have the same feed-forward response properties. Third, V1 receives afferent input from the RGC, and also contains lateral excitatory and inhibitory connections. V1's afferent and lateral connections are modified during training by divisively normalized Hebbian learning. V1 simulates roughly $3\,\mathrm{mm} \times 3\,\mathrm{mm}$ of cortical area.*
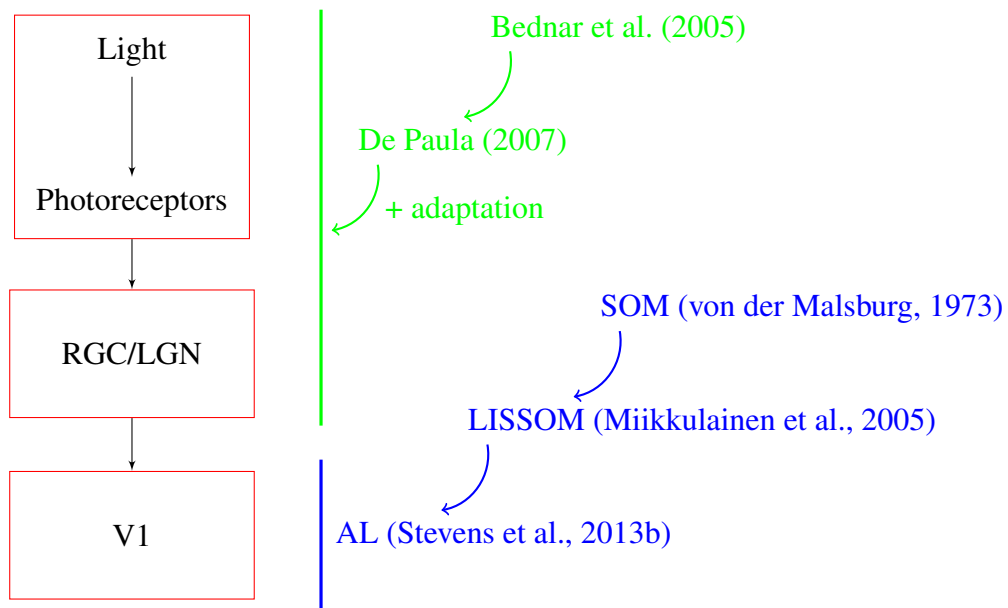
Figure 4.2: *Model's ancestry. The subcortical model is based on that of De Paula (2007), which in turn is based on that of Bednar et al. (2005) (see chapter 2); we add adaptation based on the input data. The cortical model is AL, which comes from a family of models that elaborate the self-organizing map algorithm of von der Malsburg (1973), and importantly also includes adaptation.*

There are two important issues regarding our (and previous work's) estimate of cone excitations. The first is that the LMS sensitivities we have used are corneal sensitivities, derived from human color-matching experiments. Therefore, the sensitivities do not include the effects of pre-receptoral filtering. In humans and animals, the final photopigment isomerization rate for each cone class is proportional to its corresponding corneal sensitivity, but the constant of proportionality differs between cone classes (because e.g. macular pigment absorbs more short-wavelength light than long-wavelength light). However, we will argue below that the normalization of cone sensitivities is something that must vary, so any fixed constant of proportionality between the corneal sensitivity and and photopigment isomerization can be ignored for our long-term development model. Some image databases do include an estimate of the photopigment isomerization rates; for the phenomenon we are modeling (the long-term development of color selectivity), we obtained qualitatively similar results from training on a subset of one of these (UPenn Calibrated Image Database (UPCID; Tkačik et al., 2011)).

The second issue regarding the estimation of cone excitations is the accuracy of conversion of scene radiances, as captured by an RGB camera with sensitivities that differ beyond a linear transform from human LMS sensitivities, to LMS cone excitations. As noted in chapter 3, more accurate conversions of scene radiances to cone excita-

(a) Image

(b) L receptors   (c) M receptors

(d) L-M On and Off RGC/LGN   (e) M-L On and Off RGC/LGN

(f) V1: initial   (g) settling step 1   (h) step 6   (i) step 11   (j) step 16
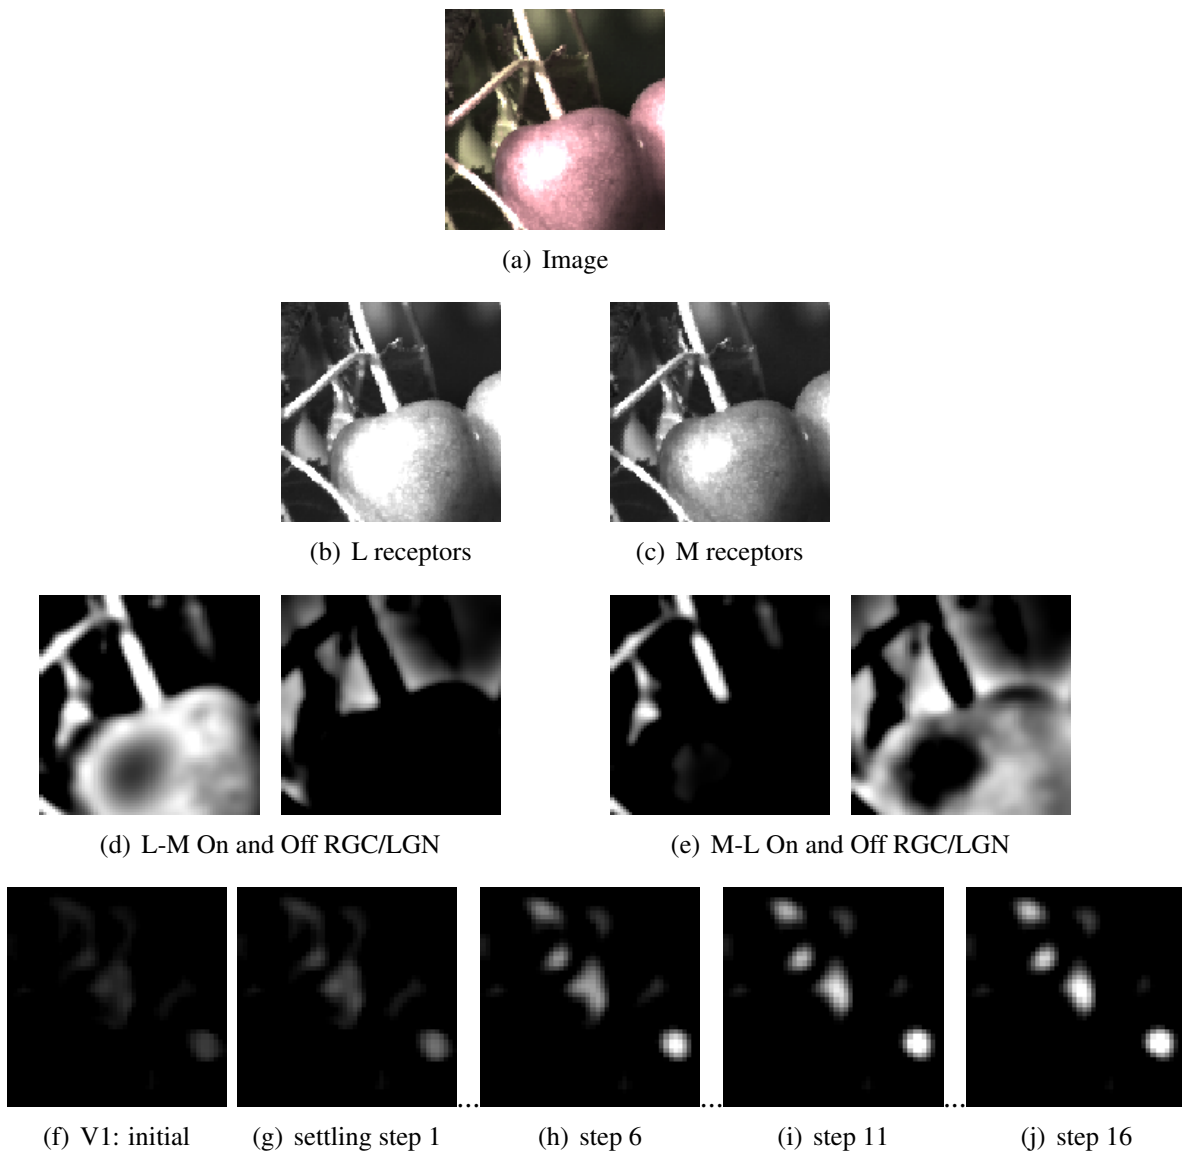
Figure 4.3: *Demonstration of one iteration in the model during training. Activities are shown for a subset of the model from figure 4.1(b). (a) At the start of each iteration, a random patch (at a random zoom level and orientation) of a randomly selected calibrated color image is presented to the model (uncalibrated RGB shown here). The processing that follows represents one visual fixation. (b) and (c) L and M cone photoreceptor responses (S not shown). (d) and (e) L-M On/Off and M-L On/Off responses (LUM and S/(L+M) not shown). (f) Initial diffuse V1 response. (g)–(j) V1's lateral connections settle the activity into discrete bubbles (multiple "winning" patches). Connections to the neurons remaining active strengthen via Hebbian learning, causing local patches of V1 to become selective for particular input patterns. The short-range lateral excitatory connections cause nearby neurons to have similar responses, while the long-range lateral inhibitory connections cause more distant neurons to learn different inputs.*
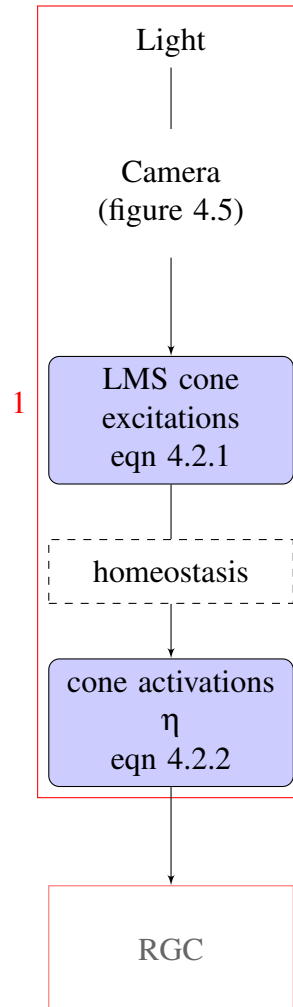
Figure 4.4: *Naive model's photoreceptor processing. Cone excitations are linearly related to photon absorption, but subsequent cone activations (responses) are not linear. The naive model includes simulated adaptive homeostasis.*
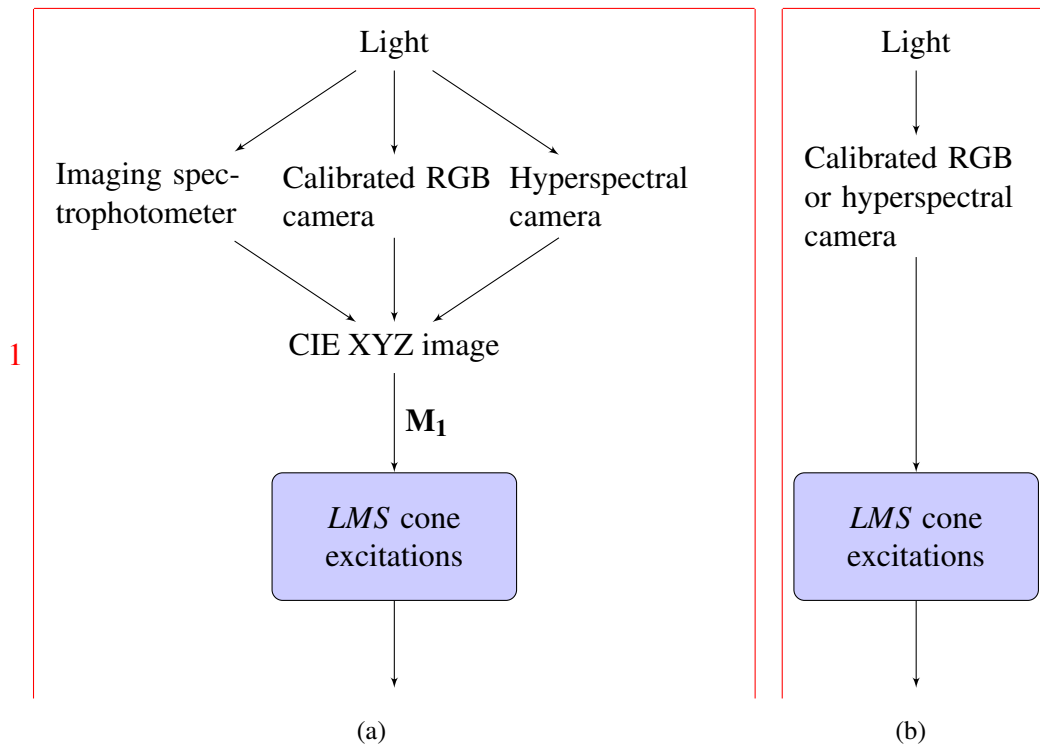
Figure 4.5: *Options for supplying cone excitations. The model can work either with images specified in CIEXYZ, or directly with human LMS cone estimates. (a) Typically, CIEXYZ is used, because it is a standard format for device-independent color, and most calibrated color image data is available in it, or can be converted to it. The CIEXYZ cmf are approximately related to LMS sensitivities by a linear transform. (b) Alternatively, if cone excitations have already been estimated, these can be used directly. This path may be used to allow external computation of LMS excitations, for instance if a higher accuracy is required (see chapter 3).*

tions are possible (e.g. by capturing images with a hyperspectral camera, or by using a more sophisticated RGB camera-to-LMS conversion). Therefore, the model allows LMS cone excitations to be supplied directly instead (figure 4.5). For the phenomena we are modeling, we obtained qualitatively similar results using Párraga et al. (2010)'s more accurate LMS cone excitation estimations. We expect the same would be true using hyperspectral images converted to cone excitations. As pointed out in chapter 2, pre-receptoral filtering varies between individuals, and over one individual's lifetime, so it is not clear that highly accurate estimation of LMS excitations is meaningful for our modeling, and requiring it would greatly reduce the range of image databases with which the model can be tested.

4.2.1.1 Adaptive homeostasis

M_1 (or any other conversion method) includes an assumption about how to normalize the LMS cone responses (chapter 3). The De Paula (2007) model uses LMS cone activations computed by Olmos and Kingdom (2004). The Olmos and Kingdom (2004) conversion from RGB camera responses to LMS responses assumes equal activity of L, M, and S cones for equal-energy (physically white) light (along with some other constraints from psychophysics—see chapter 3). In the De Paula (2007) model, this is followed by ad hoc manual scaling of each of the LMS channels in order to achieve a CR preference map that represents all three channel types. That is, the three channel scalings are parameters the modeler sets manually via an iterative procedure based on the resulting CR and OR maps, and are applicable only to a specific set of images (problem **P7-LMSscaling**).

The requirement to scale each channel differently after converting the RGB camera images to LMS activations arises because of a combination of:

- an implicit assumption that each class of photoreceptor is always equally activated by equal-energy white light;

- an image database having differing total energies over the L, M, and S cone wavelength sensitivity ranges.

We saw in chapter 2 that color depends strongly on both spatial and temporal context, and that one cause of this is likely to be neural adaptation: the L, M, and S photoreceptors do not always respond equally to equal-energy light (or whatever other fixed criterion could be used to determine relative photoreceptor responses); chapter 2's figure 2.12 showed the effect of adaptation on cone responses. Therefore, we simulate

adaptive homeostasis of the photoreceptors (to be described in detail below), both reducing the number of free parameters, and making it easier to use different/larger image databases. Note that this adaptation would not necessarily have to occur in the receptors themselves, but could occur in the horizontal cells or any other part of the pathway before the cone-opponent output from the ganglion cells.

In the following chapter we will be comparing various conditions in the retina; to do so fairly it is crucial to have precisely balanced results. Therefore, we use a batch rather than heuristic method to simulate homeostasis. However, we expect similar results would be achieved with an incremental homeostatic mechanism like that used in the model's V1 (to be described in section 4.2.3).

The cone activations at each input location i are computed as:

$$\eta_i = \sigma \begin{pmatrix} g_{L,i} L_i \\ g_{M,i} M_i \\ g_{S,i} S_i \end{pmatrix} \tag{4.2.2}$$

The constant gain factors $g_{L,i}$ are pre-computed independently for each channel, ensuring that over the course of a simulation, the cumulative total excitation of each photoreceptor type at a given location i is equal:

$$g_{L,i} = \frac{\max\left(\langle L_i \rangle, \langle M_i \rangle, \langle S_i \rangle\right)}{\langle L_i \rangle} \quad g_{M,i} = \frac{\max\left(\langle L_i \rangle, \langle M_i \rangle, \langle S_i \rangle\right)}{\langle M_i \rangle} \quad g_{S,i} = \frac{\max\left(\langle L_i \rangle, \langle M_i \rangle, \langle S_i \rangle\right)}{\langle S_i \rangle}$$
$$\tag{4.2.3}$$

where $\langle \cdot \rangle$ denotes an average over time.

As mentioned above, although the gain factors computed here result in precisely balanced long-term channel activations, this precise balance is necessary only for fair comparisons of various conditions in the following chapter. If instead of computing the gain factors per receptor, we compute one gain factor per channel (i.e. to be applied to all receptors of the same type, regardless of location), we obtain similar results. This is because, over the course of the entire simulation, each receptor experiences a very similar distribution of input, so for long-term developmental simulations it is not necessary to consider the precise spatial extent of any normalization that may occur in the retina. However, as discussed in chapter 6, such effects would likely be important over short timescales.

σ is a piecewise-linear function that keeps the overall mean activity output from the retina approximately constant at each iteration, while preventing any photoreceptor activation being larger than 1. In the next chapter, we will be changing various conditions

in the retina; σ is required to ensure there are no spurious effects from changes to the overall retinal activity.

$$\sigma(x) = \begin{cases} 0 & x \leq 0 \\ \frac{x}{\theta} & 0 < x \leq \theta \\ 1 & x > \theta \end{cases} \tag{4.2.4}$$

where θ is the same for every photoreceptor, and is computed at each iteration as:

$$\theta = \frac{\overline{X}}{m_p} \tag{4.2.5}$$

where $\overline{X}$ is the average across space of all incoming activation (i.e. across all gain-adjusted initial photoreceptor activations), and m_p is the target mean photoreceptor activation. m_p is a constant parameter, fixed for all simulations in this thesis, chosen so that the photoreceptors' mean activation is approximately equal to that of the De Paula (2007) model. Photoreceptors saturate approximately 5% of the time.

Output from the photoreceptors is next processed by sheets of center-surround ganglion cells, representing the RGC and LGN , which we go on to describe in the next section.

4.2.2   Stage 2: RGC/LGN

The model's sheets of RGC/LGN cells represent all the early visual system processing between the photoreceptors and the LGN. This involves grouping the action of horizontal, bipolar, amacrine, and ganglion cells together, focusing only on the cells that drive the LGN (see chapter 2). The retinal ganglion cells are represented by eight different sheets, the ON and OFF pathways of the four typically assumed classes of P retinal ganglion cell. At each RGC spatial location, there are therefore eight retinal ganglion cells. Each retinal ganglion cell has Difference-of-Gaussians (DoG) connection fields (CFs) with fixed weights (defined below) from various combinations of cone cells (figure 4.7). The cone-opponent channels (e.g. S ON center, LM OFF surround) primarily carry chromatic information, while the non-cone-opponent channel carries only luminance (achromatic) information because it pools all the cone types together. The naive model's RGC and LGN initial activation stage is unchanged from the De Paula (2007) model, with the exception that here we use a center-surround arrangement for the S/LM pathway, whereas the De Paula (2007) model used a spatially coextensive arrangement. The biological situation is is not clear (see chapter 2), but the behavior being tested here is not qualitatively affected by the choice.
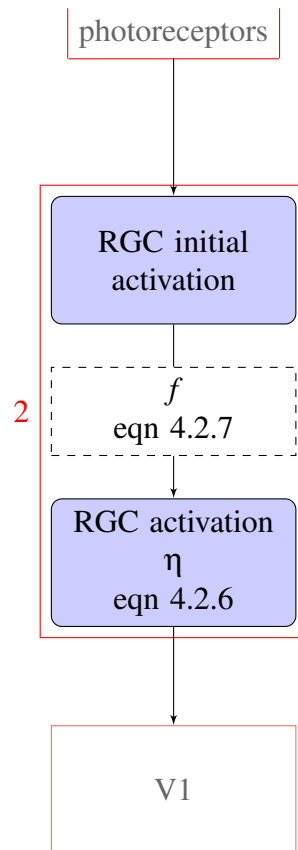
Figure 4.6: *RGC processing. Initial cone opponent activations are passed through a nonlinear function (typically half rectify).*
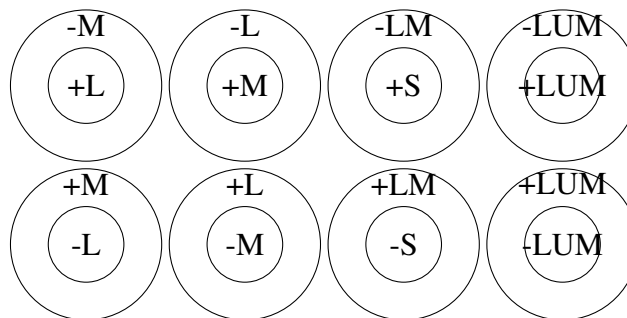
Figure 4.7: *Naive model RGC pathway. Retinal ganglion cells perform cone- and spatially opponent processing, using circular Difference-of-Gaussians connection fields. LUM is L+M+S, although some studies argue that only L and M cones are involved in LUM (see text). The "blue/yellow" (S vs. LM) pathway used in the model is center surround, although some studies argue that it may be spatially coextensive (again, see text). Neither issue matters for the phenomena being modeled in this thesis.*

In the simulations of this chapter, there are eight RGC sheets (i.e. eight classes of RGC cell: L/M On and Off, M/L On and Off, S/LM On and Off, and LUM On and Off). The activation of RGC unit j in one RGC sheet is calculated as:

$$\eta_j = f\left(\sum_\rho \gamma_\rho \sum_{i \in F_{\rho,j}} \eta_{\rho,i} w_{\rho,ij}\right) \tag{4.2.6}$$

where ρ is an index over the incoming connection types (i.e. center and surround), γ_ρ is a constant connection strength for that type of connection (inherited from the De Paula (2007) model; all connection strengths effectively equal), $\eta_{\rho,i}$ is the the activation of pre-synaptic photoreceptor i in j's CF F_ρ (i.e. the set of units to which j is connected on the input sheet), and $w_{\rho,ij}$ is the connection weight from i to j (described below).

Unlike in the De Paula (2007) model, which uses a fixed, piecewise-linear approximation to a sigmoid function, here f is a function that ensures positive activation and scales all activities to obtain a constant mean output per iteration (in a similar fashion as σ; eqn 4.2.4). The latter feature is not required in this chapter, but in the next chapter we will be manipulating aspects of the RGC/LGN and want to keep the total activation as similar as possible to avoid any spurious effects of a change in overall activity.

$$f(x) = \begin{cases} 0 & x \leq 0 \\ \frac{x}{\theta} & 0 < x \end{cases} \tag{4.2.7}$$

where θ is the same for every RGC unit, and is computed at each iteration as:

$$\theta = \frac{\overline{X'}}{m_r} \tag{4.2.8}$$

where $\overline{X'}$ is the average across space of all positive incoming activation ($X' = \max(X,0)$), and m_r is the target mean LGN activation. m_r is a constant parameter, fixed for all simulations in this thesis, chosen so that the mean RGC activation is approximately equal to that of the De Paula (2007) model.

Keeping the mean total RGC activation constant (scaling all units by the same value) is not required for the model to work, but makes comparison of the different models presented in this and the following chapter simpler. Meanwhile, half rectification (combined with V1 units having a threshold—see following section) is a reasonable model of the parvocelluar pathway being modeled here, which appears to be largely linear (Benardete et al., 1992; Benardete and Kaplan, 1997b,a; Solomon et al., 2010) (unlike the magnocellular pathway, which exhibits e.g. contrast gain control).

As mentioned above, the fixed weights $w_{\rho,ij}$ from photoreceptor i to RGC unit j are defined with a DoG kernel. The CFs for ON units have a positive center and a negative surround, while OFF units have a negative center and a positive surround. The weights $w_{\rho,ij}$ to an ON-center cell j at sheet coordinates (0,0) on one RGC sheet from sheet-coordinate location (x,y) on the photoreceptor sheet are given by:

$$w_{\rho,ij} = \frac{1}{Z_c}\exp\left(-\frac{x^2+y^2}{2\sigma_C{}^2}\right) - \frac{1}{Z_S}\exp\left(-\frac{x^2+y^2}{2\sigma_S{}^2}\right) \qquad (4.2.9)$$

The center Gaussian's width is determined by σ_C, and the surround's by σ_S (both with maximum radius R). The Gaussian widths (and maximum radii) are unchanged from the De Paula (2007) model, with a center:surround ratio of 1:4. Changing the precise ratio does not qualitatively affect the results, but the sizes may affect the degree of spatial correlation in the outgoing signal to V1 (see section 6.4.2). Note that in biology, a range of sizes, aspect ratios, and surround ratios has been observed (e.g. Tavazoie and Reid, 2000). Z_c and Z_S are normalization constants such that the center and surround weights each sum to 1. The center of the RGC CF unit j is mapped to the corresponding location on the photoreceptor sheet, so the projection is retinotopic.

4.2.3 Stage 3: V1

The model of V1 is the AL model (from Stevens et al., 2013b). The AL model is simpler than the LISSOM model used by De Paula (2007) because it requires far fewer parameters to be set by the modeler. Additionally, the AL model was found to perform better than the LISSOM model in a model of luminance (LUM)-only processing, where it developed more selective, stable, and higher quality maps, better matching animal results (Stevens et al., 2013b). Both models are elaborations of a self-organizing map algorithm introduced by von der Malsburg (1973), and allow realistic input (i.e. images rather than only abstract input), inclusion of biologically plausible properties of single neurons, and incremental learning through a biologically plausible mechanism (Hebbian learning). More background on this family of models is given in Miikkulainen et al. (2005).

Parameters are as described in Stevens et al. (2013b), where it is noted that two significant digits of precision is enough to develop qualitatively indistinguishable maps, and it is estimated that all parameters may be changed by around 10% without affecting the overall behavior.

Activation of V1 unit q is calculated as:

$$\eta_q(t + \delta t) = f\left(\sum_\rho \gamma_\rho \sum_{p \in F_{\rho,q}} \eta_{\rho,p}(t) w_{\rho,pq}\right) \qquad (4.2.10)$$

ρ is an index over the connection types: afferent ($\rho = A$; L-M On/Off, M-L On/Off, S-LM On/Off, LUM On/Off), lateral excitatory ($\rho = E$), and lateral inhibitory ($\rho = I$). γ_ρ is a constant strength for each connection type, chosen such that activity bubbles form, by balancing both inhibition vs. excitation and afferent vs. lateral activation. $\eta_{\rho,p}$ is the activation of pre-synaptic unit p in q's CF F_ρ, and $w_{\rho,pq}$ is the connection weight from p to q (described below). After an initial update from the RGC, afferent activity remains constant, while V1 activity settles over 16 steps through activation of the short-range excitatory and long-range inhibitory lateral connections. The number of settling steps was constant, and determined in advance through observation: after 16 steps, V1 activity no longer changed. f is a threshold function that ensures positive activation, but also implements single-neuron homeostasis by varying the threshold (figure 4.8), as described in the next section.

4.2.3.1 Adaptive homeostasis

While the developmental path of CR maps is not known, other features (e.g. OR) are known to develop in a stable way in animals (meaning the earliest measurable maps are similar to the eventual adult maps; see review in Stevens et al. 2013b). Therefore, it seems reasonable to assume the same for CR. However, the model of V1 used by De Paula 2007 (LISSOM) does not develop in a stable way (Stevens et al., 2013b). LISSOM's V1 activation level is maintained by periodic manual parameter adjustments, which leads to the CR map state changing suddenly at these adjustment points (and requires manual setting of various parameters for different datasets). Therefore, we instead use the AL model from Stevens et al. (2013b), which has an independent adaptive threshold θ for each neuron, automatically adjusted to maintain a fixed target activity (figure 4.8).

After the final response is calculated following settling (equation 4.2.10), the threshold of f is updated. To set the threshold for activation, each V1 unit q calculates a smoothed exponential average of its settled activity:

$$\overline{\eta_q}(t) = (1 - \beta)\eta_q(t) + \beta\overline{\eta_q}(t - 1) \qquad (4.2.11)$$

where β is a parameter determining the degree of smoothing in calculating the average.
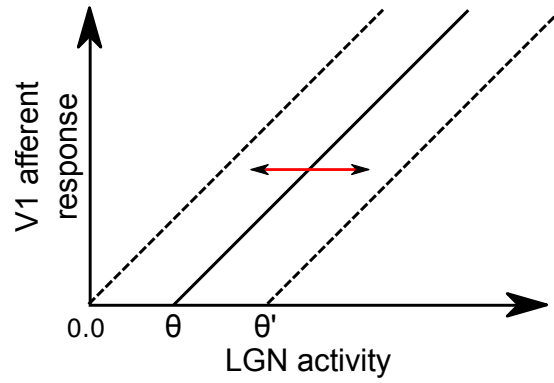
Figure 4.8: *V1 homeostasis. Each V1 neuron has an independent adaptive threshold, automatically adjusted to maintain a target activity. Image from Stevens et al. (2013b).*

$\overline{\eta_q}$ is initialized to the target average V1 unit activity μ ($\overline{\eta_q}(0) = \mu$). The threshold update is given by:

$$\theta_q(t) = \theta_q(t-1) + \lambda \left(\overline{\eta_q}(t) - \mu \right) \tag{4.2.12}$$

This brings the average activity of each V1 unit closer to the specified target.

4.2.3.2 Learning

Afferent and lateral inhibitory weights to V1 neurons are initially random, scaled with a 2D Gaussian profile, and learn during training iterations, while lateral excitatory weights are fixed 2D Gaussians. Disabling learning of lateral excitatory connections speeds up simulations, and for LUM-only models has been found to have no visually distinguishable effect on map development (Stevens et al., 2013b).

For V1 unit q at sheet-coordinate location (0,0), the initial weights from pre-synaptic unit p in connection ρ are given by:

$$w^0_{\rho,pq} = \frac{1}{Z_{\rho,q}} u \exp \left(-\frac{x^2 + y^2}{2\sigma_\rho{}^2} \right) \tag{4.2.13}$$

where (x,y) is the sheet-coordinate location of the pre-synaptic neuron, $u = 1$ for the lateral excitatory projection ($\rho = E$) or is drawn from a uniform random distribution on [0,1] for the afferent ($\rho = A$) and lateral inhibitory ($\rho = I$) connections, σ_ρ determines the width of the Gaussian for each connection type ρ, and $Z_{\rho,q}$ is a constant to normalize the total of all weights $w_{\rho,pq}$ to neuron q of type ρ to 1.0. Each connection type's weights have a maximum radius r_ρ, outside of which they are set to 0.

During training, after the settled response is calculated (equation 4.2.10), the afferent

($\rho = A$) and lateral inhibitory ($\rho = I$) weights are adjusted by unsupervised Hebbian learning with divisive normalization:

$$w_{\rho,pq}(t) = \frac{w_{\rho,pq}(t-1) + \alpha_\rho \eta_q \eta_p}{\sum_u (w_{\rho,uq}(t-1) + \alpha_\rho \eta_q \eta_u)} \tag{4.2.14}$$

where $w_{\rho,pq}$ is the connection weight from pre-synaptic unit p to post-synaptic unit q, α_ρ is the learning rate for the connection type (afferent α_A, lateral inhibitory α_I), η_q is the settled activity of unit q, and η_p is the activity of pre-synaptic unit p. For afferent connections, u in the denominator is over all afferent connection fields, i.e. all afferent connections are jointly normalized. The divisive normalization used here prevents unconstrained growth of the weights, and stands in for a biologically plausible mechanism such as multiplicative synaptic scaling or a sliding threshold for plasticity (Miikkulainen et al., 2005).

4.3 Results

Having extended the De Paula (2007) model by adding homeostatic adaptation at the photoreceptor and V1 stages, we now go on to evaluate this new model (the naive model) against the experimental data. The new model should perform at least as well as the De Paula (2007) model, even though it now has many fewer parameters (in particular, avoiding three arbitrary scaling factors for the LMS channels, and over 40 V1 scheduled parameter changes during the simulation). Along the way, we will introduce new analyses to allow close comparison with experimental data. However, we start by verifying that the two homeostasis extensions have indeed solved the previously mentioned *modeling problems*.

4.3.1 Homeostasis allows robust and stable map development

In section 4.1, we reviewed how the De Paula (2007) model requires ad hoc adjustment of relative photoreceptor channel strengths for a particular set of input images, and that the setting is chosen manually based on the appearance of the resulting OR and CR preference maps via an iterative procedure (a problem likely to exist for all previous models; problem **P7-LMSscaling**). Apart from giving the modeler extra parameters to tweak to match experimental results, the practical consequence of this is that we cannot easily use different selections of images because of the required parameter search every time.

A second problem with existing models that we reviewed in section 4.1 is that map development does not proceed in a stable way (problem **P8-V1StableDev**). In animals, we know that the selectivity of an OR preference map increases at a similar rate to its stability, whereas the model of V1 used by De Paula (2007) develops selectivity far more rapidly than stability (Stevens et al., 2013b). Apart from not matching experimental data, unstable map development makes it difficult to investigate the representation of CR, because CR preference can change suddenly for spurious reasons. Stevens et al. (2013b) introduced a new model of V1, the AL model, in which OR maps *do* develop in a realistic and stable way (in a model of luminance processing alone).

In section 4.2, we therefore further developed the De Paula (2007) model to include homeostasis at the cortical and photoreceptor levels. Figure 4.9 shows that these mechanisms do successfully solve problems **P7-LMSscaling** and **P8-V1StableDev**. In contrast to the De Paula (2007) model, the naive model:

- Develops realistic OR maps for different image sets. This eliminates the problem of using a specific, small, hand-picked selection of images (**P7-LMSscaling**). We can therefore use different, large sets of images without making manual selections, more accurately reflecting visual experience.

- Shows more realistic progression in OR map development over time, matching animal data. Importantly, the CR map also develops in a stable way. Although there is no experimental data about the development of CR maps against which to compare, stable CR map development makes testing which factors genuinely affect the map's development much easier (a requirement for the investigations in the following chapter).

4.3.2   Comparison with experimental data

Having made some necessary improvements—adding adaptation at the photoreceptor and V1 levels—to solve problems **P7-LMSscaling** and **P8-V1StableDev**, we now go on to compare the naive model against experimental data. As reviewed in chapter 2, evidence for organization by PCS in V1 comes from several sources. However here we focus on the results of intrinsic optical imaging (iOI) studies, because optical imaging (OI) is able to determine the responses of a large enough population of cells at a high enough resolution to allow us to draw conclusions about map-level responses, and to be behaviorally relevant. We will introduce some analyses as we go, to allow for close comparison with corresponding experimental data (addressing problems **P4-peaks** and **P5-RGBWgratings**). Note that we still expect the naive model to have the same key

PCS representation problems as the De Paula (2007) model (**P1-range**, **P2-PCS**, and **P3-pref**).

4.3.2.1 Representation of OR

To study the representation of OR, we measure orientation preference and selectivity maps. Rows 2, 3, and 4 of figure 4.10 show a comparison of the naive model's maps with experimental results. In the model, these plots are obtained by presenting monochrome (BW) sine gratings of multiple orientations. As for the modeling results reported by De Paula (2007), the OR preference and selectivity maps are consistent with experimental data: we have a spatially contiguous representation of OR, in the form of a map (**ER1-ORmap**). The model's map shares features common to experimental maps, such as pinwheels, fractures, and linear zones. The map is generally selective, but as with experimental results there are areas of relatively lower selectivity.

Figure 4.9 *(following page)*: *Homeostatic adaptation improves map development and stability. The top panel shows how the De Paula (2007) model develops over time (from left to right) for four different image sets (selections from the MCID), while the lower panel shows the naive model on the same sets. At each time point, a combined OR preference and selectivity map is shown. For the De Paula (2007) model, photoreceptor and V1 scaling parameters were set by hand using an iterative procedure (based on appearance of the map), for a manually selected set of 25 images from the MCID (McGill A). These same parameters were then used to train the model on a different manually selected set of 21 images (McGill B—the same images used in De Paula (2007)—row 2), all natural McGill images (approximately 650 images, row 3), and all images from the McGill Fruits category (approximately 50 images, row 4). The top row represents the best performance of the De Paula (2007) model, but even so selectivity increases rapidly initially, but the map organization continues to change—whereas animal data indicates that OR maps increase gradually in selectivity and stability together (Stevens et al., 2013b). However, even more problematic for modeling the development of OR and CR maps is that self-organization fails to produce realistic maps at all in most cases for a given set of parameters. On the other hand, the lower panel shows that the naive model has more stable development (selectivity increases more slowly, and stability increases together with selectivity), and forms realistic maps for all image sets (no parameters altered between simulations). Additionally, unlike the De Paula (2007) model, all four simulations also developed CR maps. Stability of OR map development is investigated in more detail in Stevens et al. (2013b), though not in a model that also includes CR (as is used here).*
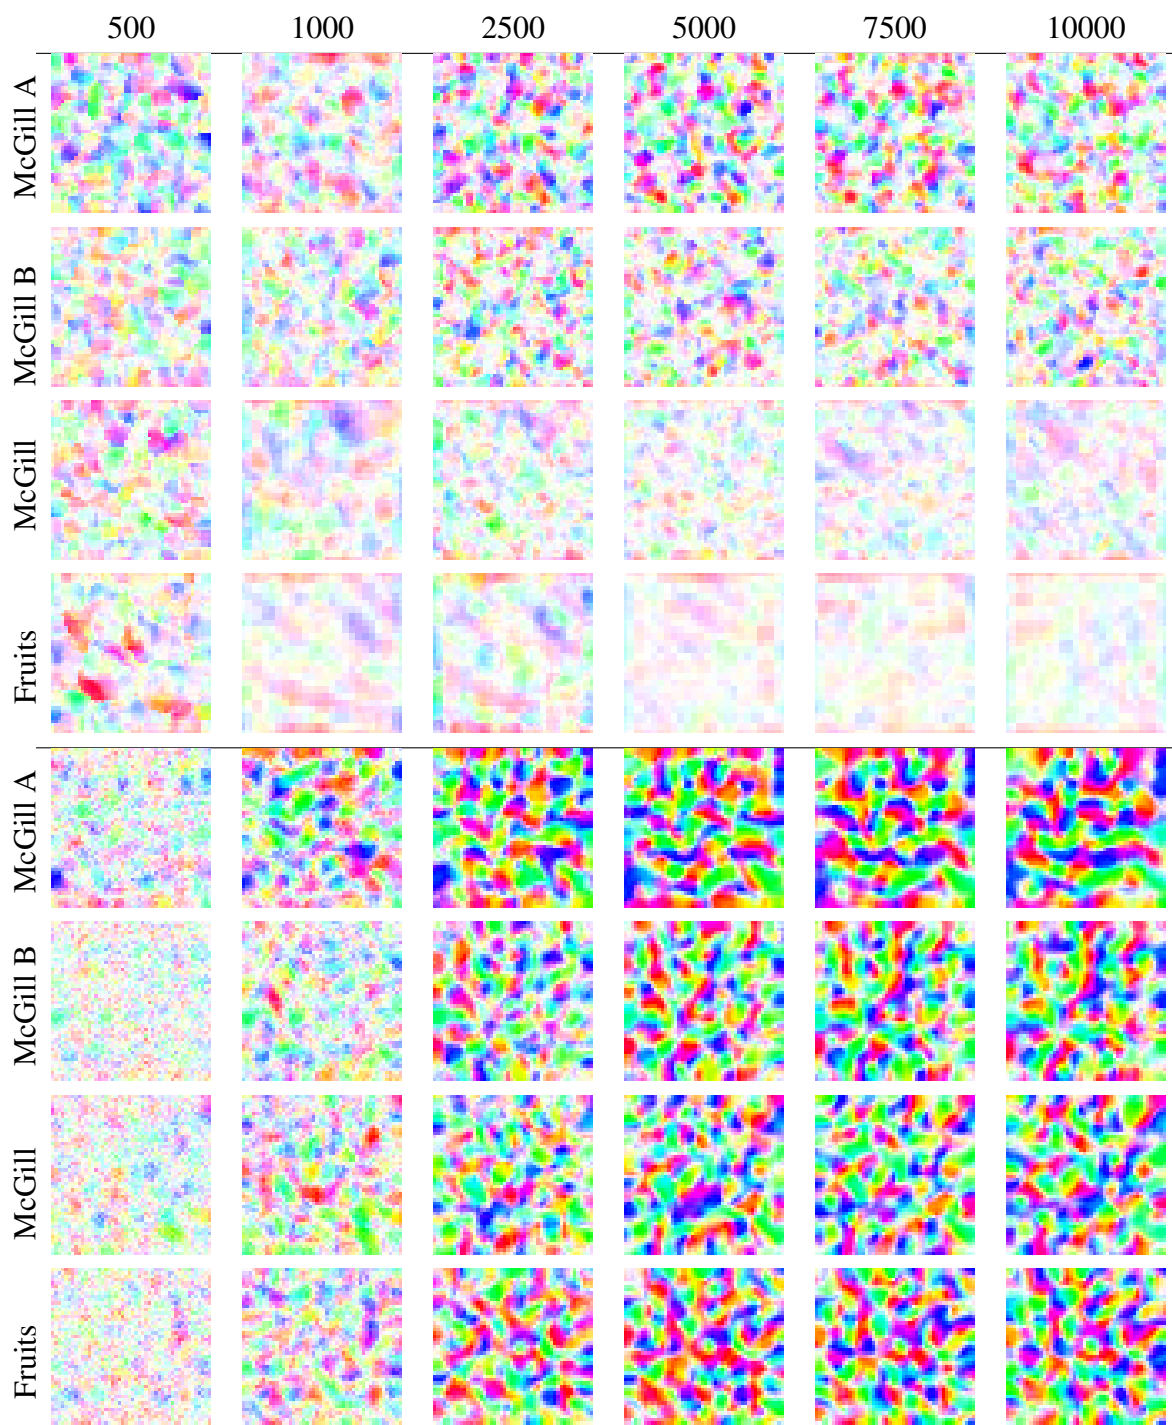
4.3.2.2 Joint representation of OR and CR

Experimental work has addressed joint OR and CR organization by comparing responses to isoluminant red-green (RG) sinusoidal gratings with the responses from black-white (BW) sinusoidal gratings of the same average luminance. The procedure for measuring such RG-BW plots in the model is similar to experimental studies (which have some differences themselves—see chapter 2).

Row 1 (RG-BW) of figure 4.10 shows that the model's joint OR and CR organization results match those from experimental studies. Note that in contrast to OR sel, black this time is the highest (difference between CR and OR), while white is the lowest. Experimental data is consistent between studies: CR patches among contiguous OR maps (**ER2-patches**). The model gives patchy responses, again consistent with the variety of experimental data. There are differences in CR patch sizes and spacing, but these are on the order of the differences between the experimental studies. Additionally, the CR patches tend to occur in areas of lower orientation selectivity, as in experimental data. We therefore draw the same conclusion as previous models (Bednar et al., 2005; De Paula, 2007), i.e. that this model matches on criteria **ER1-ORmap** and **ER2-patches**, but using measurements directly matching experimental studies (rather than only using CR selectivity).

4.3.2.3 Representation of hue within CR patches

Xiao et al. (2007) analyzed the peak responses to a selection of hues from a perceptual color space and found that those peak responses tended to fall within previously found RG-BW CR patches. Each patch responded to all or many hues (result **ER3-range**), and furthermore the distances between peaks were proporional to the distances between the hues in PCS (result **ER4-PCS**). Figure 4.11's row 1 shows the model's results compared to the experimental data. We can see that the model does not match the experiment; while the experimental data has many spatially separated hue peaks per patch, the model has many hues on top of each other. Rows 2 (CR contours) and 3 (CR peak distances) make this mismatch clearer. While in experimental data the peaks in the patches follow distances in PCS, the model has only two distance groups—there appear to be two types of patch, each representing only a limited range of hues, with many different hues not separated. As expected, then, hue representation in the naive model is still a poor match to experimental data (i.e. **P1-range**, **P2-PCS**, and **P3-pref** are still present).

What causes the naive model to fail to develop an organization of hue that matches the experimental results of Xiao et al. (2007) and Chatterjee (2010)? In the following section we hypothesize that the model lacks mechanisms to cope with variabilities to which the early visual system adapts.

Figure 4.10 *(following page)*: *Comparison of naive model and experimental data for CR vs OR organization. Experimental data was introduced in the same format in figure 2.20: each 1.5 inch square corresponds to* $3\,\mathrm{mm} \times 3\,\mathrm{mm}$ *of cortical tissue; NA indicates no data available.*

Row 1, RG-BW: black indicates greater response to red/green isoluminant gratings, while white indicates greater response to luminance (black/white) gratings. The model, like the data, shows CR patches. Patches are irregular and different sizes, as they are in experimental data. The model's patch spacing and size is consistent with experimental data.

Row 2, orientation selectivity: black (low value) indicates low selectivity (no preference for one orientation over others), while white (high value) indicates high selectivity (strong preference for one particular orientation over others). The model, like the experimental data, develops a spatially contiguous OR selectivity map. CR patches (in red, from row 1) tend to occur in areas of low OR selectivity.

Rows 3/4, orientation preference/orientation preference and selectivity: saturation represents selectivity, while hue represents preference from red (horizontal) through green (vertical) back to red (horizontal). The model, like experimental data, contains preferences to all orientations. The map is organized smoothly in general, and shares features of the experimental maps.
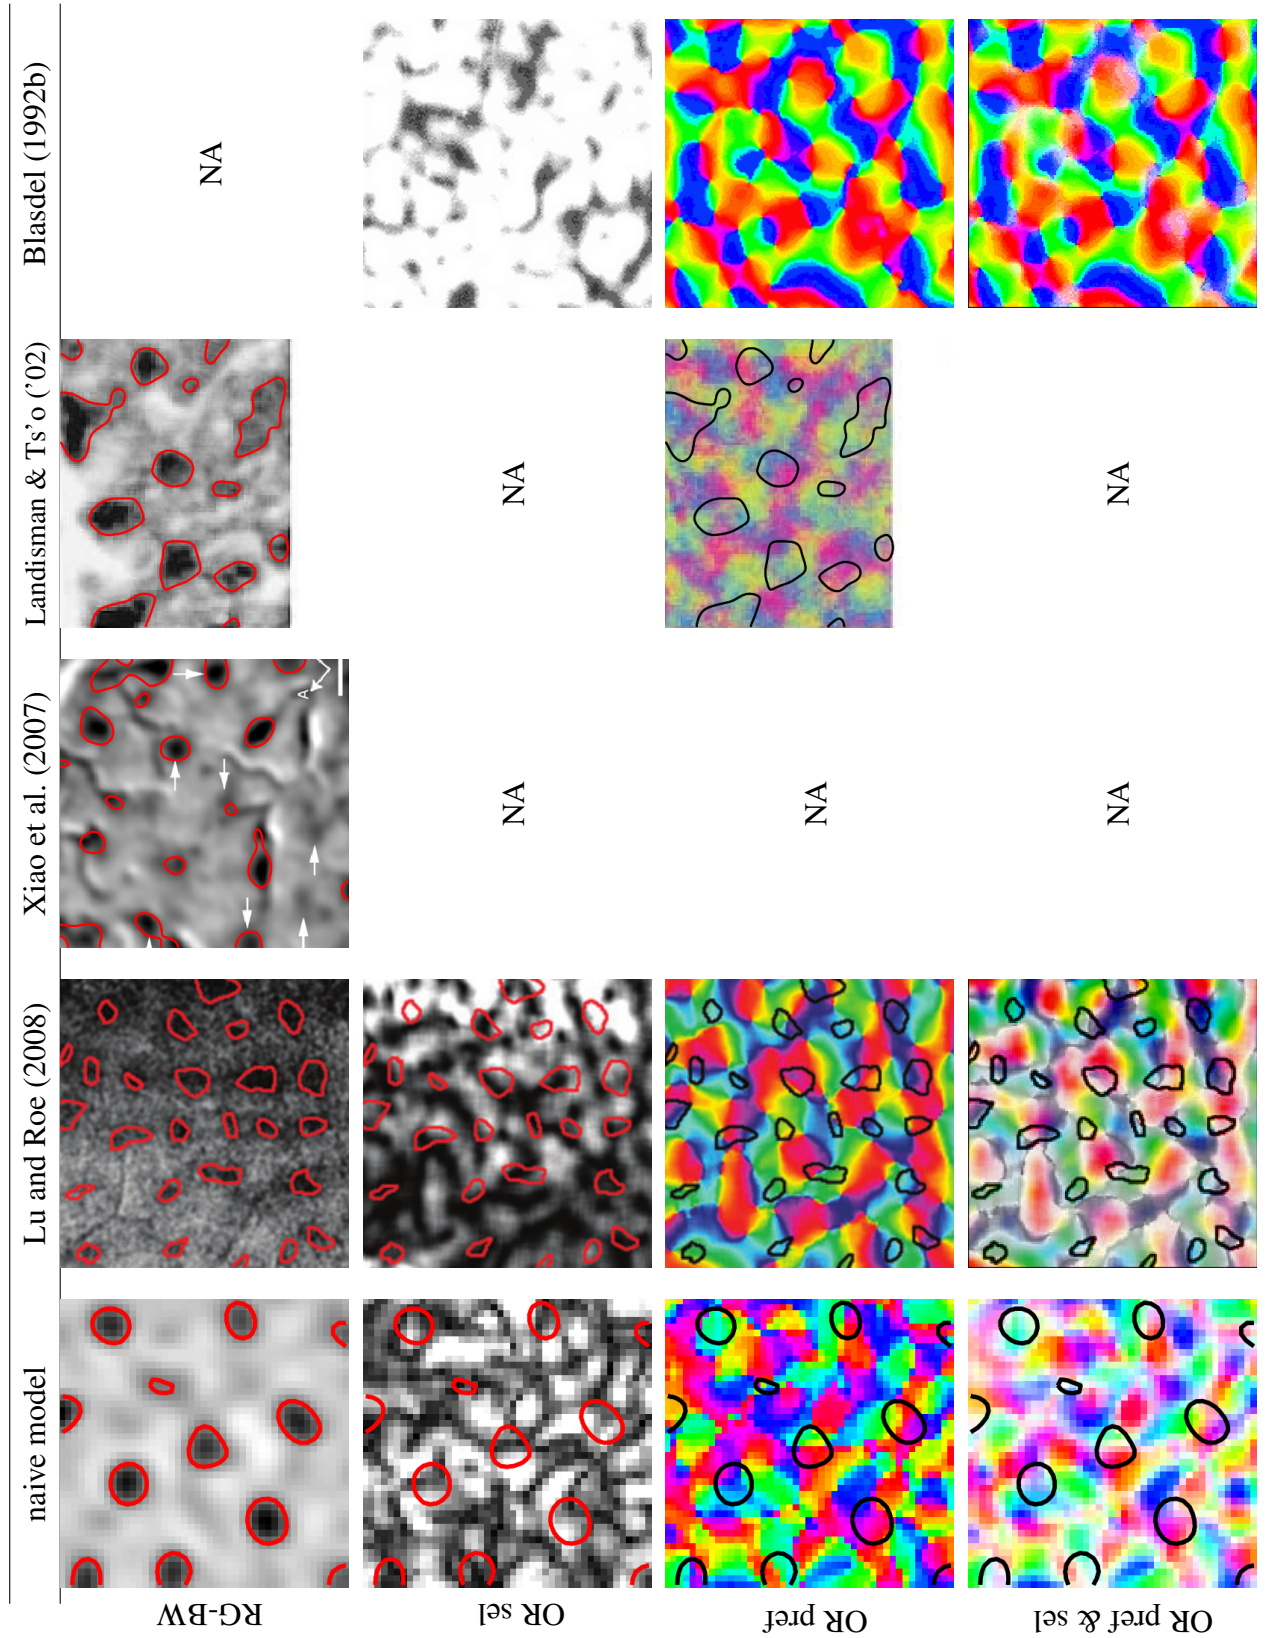
4.4 Discussion

We started with experimental results showing OR cells organized into a spatially contiguous map (**ER1-ORmap**) and CR cells in patches (**ER2-patches**). The CR patches represent a wide range or all hues (**ER3-range**), and the peak responses to various hues are spatially separated, following distances in PCS (**ER4-PCS**). We also started with the best existing model of these phenomena, the De Paula (2007) model, which develops an OR map with CR-selective cells in patches—but which does not match experimental data relating to hues within the CR patches. In order to reach a stage where we can use the model to investigate the representation of hue, we began by extending the model to solve a number of *modeling problems* (problems **P7-LMSscaling** and **P8-V1StableDev**) and to address *not directly comparable measurements* (problems **P4-peaks** and **P5-RGBWgratings**). We did this by adding adaptive homeostasis at the photoreceptor and V1 levels, and performing a detailed comparison with ex-

---

Figure 4.11 *(following page)*: *Comparison of naive model and experimental data for organization of hue. In the naive model, unlike in the experimental data, hue responses are not spatially well separated, and CR patches do not respond to all hues. Experimental data was introduced in this format in figure 2.23: each 1.5 inch square is* $3\,\mathrm{mm} \times 3\,\mathrm{mm}$ *of cortex.*
First row, CR peaks: in the model, patches of high CR selectivity are filled in white. These correspond to the RG-BW regions in figure 4.10, except they capture all color selectivity rather than only red/green. Within each patch, the peak response to each of the HSV hues presented is marked with the stimulus color (peaks can be on top of each other). In the experimental data, regions are defined (but not shown on the peaks plot) by being significantly activated to all hues, vs. gray. In the experimental data, peak responses to different hues are spatially separated, and distances correlate with distances in CIELUV. In the model, many peaks are on top of each other; hues are not well separated.
Second row, CR contours: contours drawn at a constant fraction of the peak. In the experimental data, contours are shown only for two sample patches.
Third row, CR peak distances: across all patches in the experimental data, the average distance between peaks of different hues is correlated with distance in PCS. In the model, there are only two clusters of peaks: one around red (H = 0) and one around green/blue (H = ±0.5). Note: the experimental plot was introduced in more detail in figure 2.21.
Last two rows: CR selectivity and preference maps—data from the model that is not yet available from animals. We can see that the model has developed two types of CR unit, one type with a preference for green (i.e. L and M cones), and the other type with a preference for blue/purple (i.e. S cones). The contours, drawn around the highest responses to RG-BW gratings (from figure 4.10), coincide with the green-preferring units.
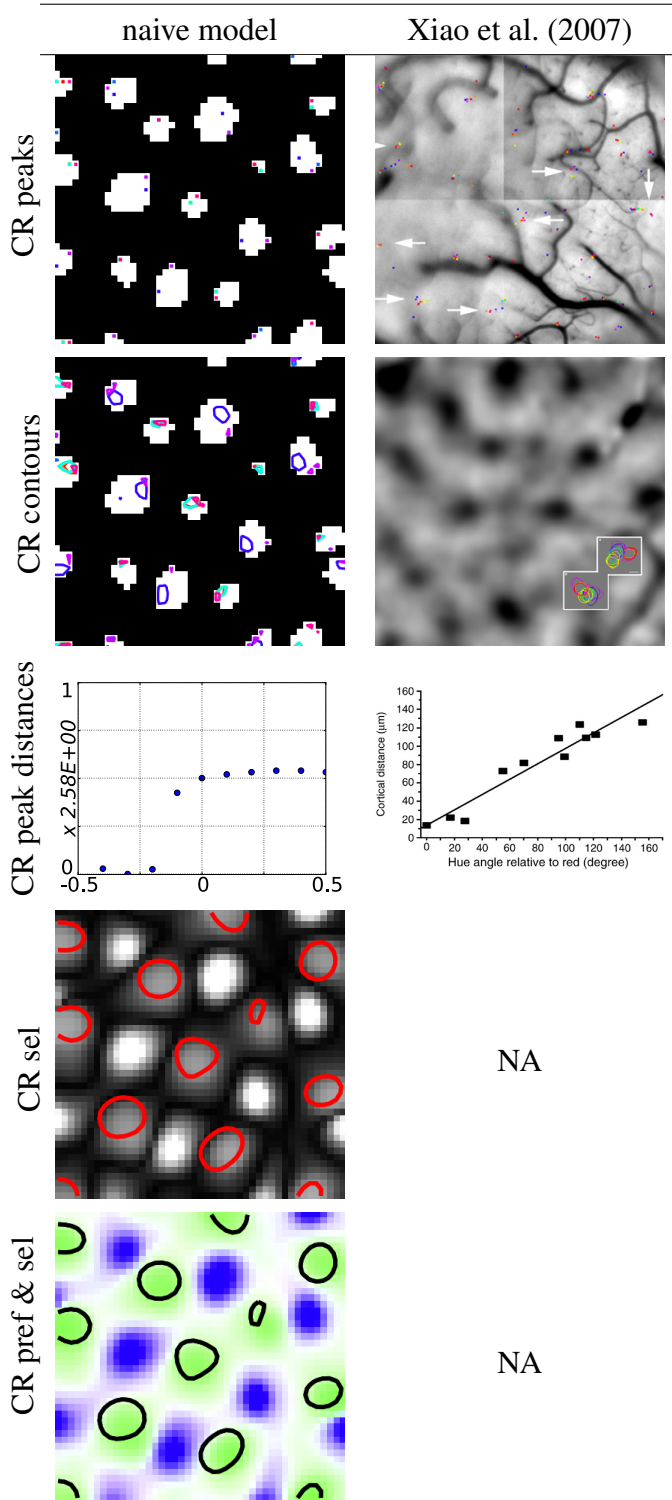
perimental data. This left us with the naive model. The organization of OR is an improvement over the De Paula (2007) model, and the joint OR and CR organization organization remains realistic. The model can now work with various datasets, with no need for images to be selected specifically, or requiring the modeler to set parameters (i.e. there are now fewer degrees of freedom). However, the naive model still does not allow coding of CR by place or preference: it retains previous work's CR representation problems: problems **P1-range**, **P2-PCS**, and **P3-pref**. The model is therefore not yet suitable as a model of CR representation in the cortex.

What causes the naive model to fail to represent hue as is found experimentally in macaque V1? The model appears to be sensitive to a number of dimensions to which the brain is relatively insensitive. As discussed in section 4.1, we know the brain is subject to a large amount of physiological and environmental variability, yet still develops a fairly consistent PCS—even though we know that visual experience plays a major role in the development of color vision (chapter 2). As an example, the model is likely to be sensitive to the highly biased hue distribution of the image databases being used (chapter 3). Competitive Hebbian learning likely amplifies the bias, and could result in a final CR map that fails to represent relatively rare hues. While human PCS may vary slightly between individuals, it is broadly consistent (chapter 2) despite natural scenes containing a biased distribution of perceptual colors (chapter 3). We think that the naive model's failures therefore likely point towards dimensions the brain is relatively insensitive to. In the following chapter, we will go on to explore these dimensions to test this hypothesis.

4.5   Conclusion

We have extended the best previous model of the development of joint OR and CR organization (De Paula, 2007), adding photoreceptor and V1 homeostasis, and V1 analyses matching experimental analyses, thus creating what we call the naive model. This model allowed us to begin investigating the representation of hue in V1. The naive model continues to match experimental data for OR maps and joint OR and CR organization (having contiguous OR maps that are better than those of the De Paula (2007) model, and still with CR patches), despite having fewer degrees of freedom. However, the naive model—like previous models—does not match experimental data about the organization of hue: while experimental data shows many or all hues are represented per CR patch, with peak responses to different hues spatially separated and following the order of PCS, the naive model represents a limited number of hues per patch, and

the hues do not show good spatial organization. We believe the naive model's failures point towards well-known mechanisms present in the early visual system but absent from the model that likely allow the visual system to cope with physiological and environmental variabilities and develop a reasonably consistent PCS. We will investigate how the model can develop a realistic representation of PCS in the next chapter.

Chapter 5

Idealized model of color map development

5.1 Introduction

In the previous chapter, we saw that the naive model—our first attempt to model the development of perceptual color space (PCS) in primary visual cortex (V1)—develops realistic joint orientation (OR) and color (CR) organization (experimental results **ER1-ORmap** and **ER2-patches**), but that critically, a hue code by spatial position or preference is not possible (experimental results **ER3-range** and **ER4-peaks**). The naive model was an improvement over the most successful existing model (with more stable and robust development of OR and CR maps), so there is now a clear challenge: can we further develop a model of this kind to successfully replicate the experimental hue organization results? In this chapter, we demonstrate that it is possible, and in doing so we will show important problems that the brain appears to be overcoming to construct PCS, but that the naive and other previous models are not.

Previously we have reviewed that it is not obvious how an organization for hue by PCS should arise:

- PCS differs from wavelength space (WS) in a number of ways (chapter 2);

- visual experience is a highly non-uniform sampling of PCS (chapter 3);

- different individuals have varied visual experiences and physiologies (chapter 2);

- an individual's physiology and experiences change over time (chapter 2).

Despite this, PCS is reasonably consistent between individuals, and over time for one

individual. On the other hand, we saw that existing models of the early visual system are unable to account for the development of PCS. Focusing on the naive model, we hypothesized that it is sensitive to a number of sources of variety that the early visual system is surprisingly insensitive to. For instance, we know that rearing environment is critical for development of color vision (chapter 2), yet the highly non-uniform sampling of PCS that typical visual experience constitutes does not greatly affect the PCS learned by the visual system. We expect that the early visual system deals with these sources of variety via well-established biological mechanisms such as adaptation and normalization—mechanisms that are, however, missing from models, in part because there is a multitude of such mechanisms in the nervous system (reviewed in chapter 2), and it has been unclear which, if any, might be relevant to color vision.

To determine which such mechanisms are missing from these models, in this chapter we will explore what modifications to the naive model are necessary for it to match experimental data. The modifications, which we refer to as *idealizations*, imply problems the visual system is overcoming in creating PCS. Each idealization is a modification along a particular *CR-biasing dimension* of the model. These dimensions of the model are dimensions that the visual system appears not to be sensitive to, but that do affect the model's CR representation (e.g. the balance of CR vs. luminance information, or the representation of different hues). Analogous to complex cells (which are relatively insensitive to position) or cells with contrast-invariant tuning, these idealizations suggest ways in which neurons are relatively insensitive to sources of variation in the environment, or to certain aspects of physiology, such as retinal cone ratios.

In the following chapter, we will relate the idealizations to biologically well established mechanisms such as adaptation and homeostasis. Such mechanisms are typically omitted from models of luminance processing because they add complexity, and the models work well enough without them. Therefore, the first goal of the current chapter is to demonstrate that at least for CR such mechanisms (which we call *rescues*) are critical and cannot safely be omitted. We will do this by creating the idealized model, containing all the idealizations together, and showing that it is able to develop a realistic representation of hue. The second goal of this chapter is then to show what effect the rescues must have on the model to achieve realistic results. We will do this by showing the range of behaviors possible from the model, delineating each dimension. This analysis will guide future implementation of rescue mechanisms. Additionally, seeing the range of behaviors possible from the model will help to interpret existing and future experimental work. So far, experimental results are limited to a few individual macaques; results may differ in ways shown by the model for different species, or for individuals reared in different environments.

5.1.1   Contributions of this chapter

1. This chapter begins by identifying the aspects of the naive model that affect CR representation (section 5.2), i.e. aspects that affect the hue representation, or the balance of OR vs. CR information. By identifying these CR-biasing dimensions in the model, we enumerate the specific problems that an organism might be addressing that lead to the development of PCS.

2. For a model with a certain position on each dimension, we show in section 5.4 that a realistic organization of OR (a spatially contiguous OR preference map) can develop alongside CR patches (figure 5.6), and that the range and spatial organization of hues represented in the patches (figure 5.7) matches experimental results from macaque monkey. This model, which we refer to as the idealized model, demonstrates for the first time that a developmental model is able to match experimental results **ER1-ORmap**, **ER2-patches**, **ER3-range**, and **ER4-peaks** (i.e. with peak hue location separations proportional to perceptual separations, within a patch, thereby allowing a code for color based on peak location or preference).

3. After demonstrating how the model can match experimental results, in section 5.5 we then analyze the model in detail, showing the range of results the model can produce, depending on architecture and parameters. This (a) demonstrates what a rescue for each dimension must achieve, and (b) shows that alternative possible organizations of CR and OR are possible, depending on physiological and environmental factors, acting as detailed predictions for future experimental work. Overall, this analysis will help us understand the space of possible mechanisms and organizations in which color vision evolved.

5.2   Background

The previous chapter's model, the naive model, matched **ER1-ORmap** and **ER2-patches** with fewer qualifications than the previous best model of the development of joint OR and CR organization (chapter 2). Critically, however, the naive model does not provide a representation of CR that could be used for hue coding by preference or spatial location. We identified the following three problems:

- **P1-range**: Each CR patch only responds to a limited range of hues around one preference (e.g. red and yellow, or blue and violet), in contrast to **ER3-range**,

which shows patches responding to many/all hues.

- **P2-PCS**: Hues are not spatially organized according to PCS. For instance, the contour plot of figure 4.11 shows largely concentric responses to different hues. This contrasts with **ER4-PCS**: within CR patches, the distances between peak responses to different hues are correlated with the perceptual distances (also meaning red and violet are closer than red or violet is to green i.e. joining of opposite ends of wavelength spectrum).

- **P3-pref**: Neurons in the model have a restricted set of hue preferences. There is no experimental data on this yet, but it shows clearly that the model's hue representation does not reflect or support PCS.

Hence, no previous modeling work has shown how a realistic range of hues can be represented within CR patches (i.e. **ER3-range**), or how a realistic spatial organization of those hues could develop (i.e. **ER4-PCS**). Therefore, as stated in the previous section, in this chapter we will extend the naive model to show how experimental results **ER3-range** and **ER4-peaks** can be matched. To do this, we must first consider what causes the naive model to fail to develop a realistic organization of hue.

What components of the model can affect CR representation, e.g. by introducing a bias in the hue representation, or by altering the balance of OR vs. CR information? That is, what are the CR-biasing dimensions of the model? The model V1 contains nothing specific to CR (i.e. it does not treat CR information any differently from luminance information, and it does not treat any hue specially), so it cannot contain a CR-biasing dimension (although its Hebbian learning mechanism could exaggerate any existing, incoming bias). However, the subcortical pathway and input data could introduce various CR biases, and any such CR bias will affect what V1 learns. Considering the subcortical pathway and input data, there are only five possible sources of CR bias in the naive model; each has a link to a source of variability that the brain may be overcoming in developing PCS:

1. Input hue bias: We saw in chapter 3 that available image databases of natural scenes have highly biased distributions of perceptual hues e.g. purple almost never occurs in natural images. Previous work examining natural scenes agrees, also indicating there is a biased distribution of colors in natural scenes (again, covered in chapter 3). The model of V1 used in our simulations uses competitive Hebbian learning, which will learn (and could exaggerate) a bias in the input. Previous work has shown that when a model representing a feature (such as spatial frequency) is given only a limited part of the range during training, the

result is an even more limited representation of that input feature (Palmer, 2009). In the case of hue, the heavily biased input datasets are a likely cause of the limited representation of color in V1. Rare colors may not be represented at all in the final map.

2. Photoreceptor correlation: The model's L and M cones are highly correlated, as for humans and macaques. Hebbian learning could, given such correlations, fail to distinguish L and M responses. This would lead to restricted hue preferences, and no spatial separation between peak responses to different hues.

   De Paula (2007) showed that a Hebbian neuron is unable to distinguish highly correlated photoreceptor channels, but this analysis did not include a model of the RGC and LGN, so we remain uncertain of the impact of photoreceptor correlation on the naive model. We expect that the cone-opponent retinal ganglion cells should reduce correlation (chapter 2); the L-M and M-L RGC channels should be less correlated than the L and M photoreceptor channels. However, there are still spatial correlations revealed by the spatially center-surround nature of L-M and M-L cells. The model V1 could therefore still be unable to distinguish the L-M and M-L channels, depending on their level of correlation in a particular dataset.

3. Relative photoreceptor channel strengths: A particular database of images may have more energy at e.g. L cone wavelengths than at S cone wavelengths (chapter 3), so there may be differing overall activities in the three photoreceptor classes over the course of the simulation. We already know this has an effect on CR representation; we solved **P7-LMSscaling** in the previous chapter by simulating photoreceptor adaptive homeostasis to balance the overall channel activities. The energy absorbed by the different photoreceptor classes over any period of time (short or long) is unlikely to be equal. Even if the average spectral power distribution (SPD) coming into the eye were flat (white light), at any particular small location on the retina an L cone would absorb more energy than an adjacent M or S cone (the lens and macular pigment filter short-wavelength light more than long-wavelength light, and also the L cone has the widest sensitivity of the three cone types). We did not simulate uneven ratios or distributions of L, M, and S cones (as found in biology—chapter 2), but uneven ratios globally or locally would additionally affect this.

4. RGC pathway symmetry. The naive model does not have a symmetric RGC pathway: the RGC and LGN cone-opponent pathway does not oppose all cone types. For example, there are L-M and M-L channels, but no L-S or S-L. We

do not know the effect of the unsymmetric pathway on hue representation in the model—the pathway may cause certain hues to be de-emphasized, for instance. In macaque, the retinal ganglion cells do generally appear to cluster into these restricted opponencies (chapter 2), and there may be a good explanation for the RGC's unsymmetric pathway; e.g. from the perspective of information processing, high L and M correlation means that having both S/M and S/L channels would be redundant (chapter 2), or from the perspective of developmental mechanisms, random wiring with few S cones may lead primarily to L vs. M retinal ganglion cells (also chapter 2). The implications of these asymmetries for development and function are not yet clear, because no model has compared symmetric vs. non-symmetric pathways.

5. Luminance vs. CR: Apart from aspects of the model likely to specifically affect the representation of hue, we also consider an aspect of the model that could affect overall CR organization: the weighting of the luminance and CR RGC and LGN channels. In the model, each V1 unit receives input from both luminance and CR channels; competitive Hebbian learning will result in the feature causing most variance to dominate (De Paula, 2007). To get patches of CR with many hues represented in one patch, alongside a largely contiguous and generally selective OR map, the balance of luminance and CR inputs must be such that OR dominates. If it were the other way round in this type of Hebbian-learning model, CR would dominate and would not be constrained to patches; we would not get all hues within a CR patch, but instead would have a large-scale organization into separate color patches. For instance, Miikkulainen et al. (2005) showed that the balance of OR and motion direction (DR) preference maps can be altered by adjusting input pattern speed during training. At low speeds, OR dominates and is the largest scale feature after training, whereas at higher input speeds, DR dominates and becomes the largest scale feature. The right balance between luminance and CR input must therefore be necessary to match the particular macaque results, and it could be that other balances match different species and/or individuals, depending on rearing environment and physiology. This issue is relevant to the debate on the separation of OR and CR processing in V1, as introduced in chapter 2.

To test whether the above aspects do affect the representation of CR in the model's V1, and to discover if this kind of model can develop a realistic representation of hue in V1 at all, we will construct a model eliminating all these sources of CR bias—the idealized model. Building such model will require several biologically implausible manipulations, but the goal is to determine whether a model of this kind can replicate

experimental results **ER3-range** and **ER4-peaks** at all—and while doing so to understand problems the visual system is apparently overcoming in creating PCS. That is, we are removing the effects of dimensions we think the brain may be relatively insensitive to. We will see that doing this does allow the model to develop a realistic organization for hue, matching experimental results for the first time.

Each change to create the idealized model takes the model away from the previously established biologically realistic setting, and forms a dimension containing both the biologically realistic (naive model) parameters and those that give the best match to known experimental results (idealized model). In the second half of the results section, we will examine these dimensions, showing the range of behaviors possible, and thus illustrating what a rescue for each dimension would need to achieve. This will give us a better understanding of both the model and the problems the visual system is apparently overcoming.

In the following chapter, we will discuss biologically plausible rescues for each dimension, guided by the results in this chapter. However, first we will simply show realistic results are possible, what is required to achieve that, and how these changes affect the behavior of the model. The following section describes how we implement the idealizations and dimensions.

5.3 Methods

Chapter 4 introduced a model of the early visual system, the naive model, consisting of sheets of firing-rate–based point neurons representing photoreceptors, RGC, and V1. The model V1 self-organizes through Hebbian learning of its afferent and lateral connections as the simulation proceeds, and we can measure a realistic OR map with CR patches. In this section, we will extend the naive model to allow us to perform all of the idealizations introduced in the previous section—simulating the model being relatively insensitive to certain sources of variation (physiological and environmental variabilities that the visual system appears to compensate for). We will also introduce how we can explore the model's CR-biasing dimensions, to be used in the second part of the results section (5.5), by controlling the position along each dimension with a parameter. All the modifications we will add are prior to V1 (i.e. we only modify the photoreceptors and RGC), and are shown in overview in figures 5.1 (input), 5.2 (photoreceptors), and 5.4 (RGC and LGN). Parameters are listed in appendix C, and the model and simulator used to generate results in this chapter are freely available

from www.topographica.org.

5.3.1   Input: hue jitter

We begin by allowing control over input to the model. In section 5.2 we pointed out that the non-uniform sampling of PCS that natural image databases represent likely affects hue representation in the model. To test this, we can ensure the incoming hue distribution is uniform in PCS, covering all hues equally. For each image patch presentation, the source image—typically in CIE 1931 XYZ color space (CIEXYZ)—is converted to HSV color space (HSV; Smith, 1978), in which H, S, and V are independent perceptual dimensions. While HSV is not a perfect perceptual space (no space is yet—see chapter 2), this transform allows a simple way to alter the hue distribution with minimal effect on other aspects of the images. Chapter 3 showed the procedure in more detail, but briefly, across the whole patch, every pixel's *H* (which varies circularly from 0 back to 1) has a uniform random number between 0 and 1 added to it, before being converted back to its original space (e.g. CIEXYZ) for use by the simulation as before. Over 10,000 image patch presentations, the incoming distribution of hues is therefore uniform in an HSV-based approximation to PCS. In this chapter, we use HSV color space; Smith, 1978 (HSV) computed from sRGB as the perceptual space, because we will be using sRGB as our idealized receptor space (explained in the following section).

The scenario described above, with an unbiased (uniform) incoming hue distribution, is the idealized position on a dimension describing the hue bias. This dimension, which will be explored in the second part of the results section, is controlled by d_{HB}. The dataset's original hue distribution is the full bias of the dataset, $d_{\text{HB}} = 1$, meaning the hue distribution is unchanged. The full procedure is below, and is also indicated on figure 5.1:

1. Convert to sRGB:

$$\begin{pmatrix} R_i \\ G_i \\ B_i \end{pmatrix} = \mathbf{M_2} \begin{pmatrix} X_i \\ Y_i \\ Z_i \end{pmatrix} \quad (5.3.1)$$

2. Convert to HSV:

$$\begin{pmatrix} H_i \\ S_i \\ V_i \end{pmatrix} = \text{g} \begin{pmatrix} R_i \\ G_i \\ B_i \end{pmatrix} \quad (5.3.2)$$

    where g is the standard non-linear transform of RGB to HSV (Smith, 1978),
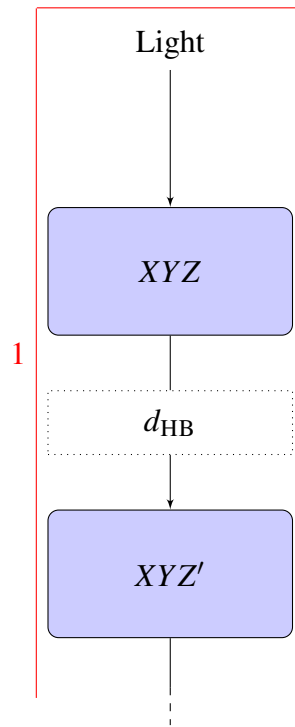
Figure 5.1: *Idealized model: input processing. The idealized model starts from an "unbiased" position, and allows CR biases to be added back in as a dimension. The input dimension manipulation is shown in a dotted box: altering the incoming distribution of hues (input hue bias; d_{HB}).*

from Python 2.6.4.

3. Jitter the hue:

$$H_i' = H_i + \mathrm{U}(0, 1 - d_{\mathrm{HB}}) \tag{5.3.3}$$

where U is a number drawn from a uniform random distribution, ranging from 0.0 to 1.0.

4. Convert back to CIEXYZ:

$$\begin{pmatrix} X_i' \\ Y_i' \\ Z_i' \end{pmatrix} = \mathbf{M_2}^{-1} \mathrm{g}^{-1} \begin{pmatrix} H_i' \\ S_i \\ V_i \end{pmatrix} \tag{5.3.4}$$

We will use the same image database as the previous chapter (BCID), but we expect that any database of natural scenes would give similar results since this procedure normalizes the hue distribution.
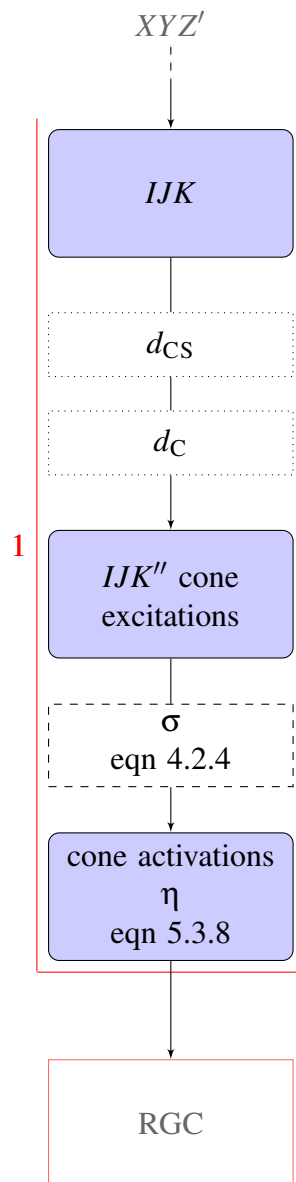
Figure 5.2: *Idealized model photoreceptor processing. The idealized model starts from an "unbiased" position, and allows CR biases to be added back in as dimensions. Dimensions are shown in dotted boxes: altering the relative excitation of the different photoreceptor classes (relative photoreceptor channel strengths; d_{CS}), and the correlation between photoreceptor classes (photoreceptor correlation; d_C). IJK'' represents this model's three cone excitations (the equivalent of the naive model's LMS—c.f. naive model overview, figure 4.4). After excitations and dimension manipulations, η are the photoreceptor responses, passed onto the retinal ganglion cells in the next stage (where further dimensions are applied).*

5.3.2 Photoreceptors

In section 5.2 we pointed out two aspects of the photoreceptors that could affect hue representation in the model: unbalanced correlation of the photoreceptors (*photoreceptor correlation*), and unbalanced activities overall in the different photoreceptor classes (*relative photoreceptor channel strengths*). We will modify the simulation to allow both of these effects to be removed (for the idealized model) and then explored. We already know the second has an effect from the previous chapter, but in this chapter we will systematically explore that effect. Figure 5.2 shows all the steps, which we now go on to describe in detail.

As seen in chapter 3, human/macaque LMS cone sensitivities are relatively broad and overlapping, leading to highly correlated absorptions. In particular, the L and M cones are much more correlated than L and S or M and S. There are also correlations from the incoming data, which can also cause correlations to vary between different cone types. To learn the effect of this on the model, we can instead use evenly spaced cone sensitivities. We could do this by using hyperspectral images, from which we could simulate any type of cone absorptions desired, but there is only very limited hyperspectral data available (chapter 3). Instead, as mentioned in the previous section, we convert the images to sRGB color space; Stokes et al., 1996 (sRGB), computed from the Barcelona Calibrated Image Database; Párraga et al., 2010 (BCID) CIEXYZ dataset. This transformation alters the cone sensitivities as shown in figure 5.3: the sensitivities are more evenly spaced, with more equal overlaps. In a subsequent section, we will add a mechanism to gradually increase correlation, allowing the effect to be explored as a dimension.

The naive model's *LMS* cone excitations (equation 4.2.1) are thus replaced by three channels we call *IJK* (which, although they begin as *RGB*, will be modified in subsequent sections—hence *IJK* rather than *RGB*):

$$\begin{pmatrix} I_i \\ J_i \\ K_i \end{pmatrix} = \mathbf{M_2} \begin{pmatrix} X_i \\ Y_i \\ Z_i \end{pmatrix} \tag{5.3.5}$$

5.3.3 Photoreceptors: relative strengths

As documented in chapter 3, the total L, M, and S excitations are not equal for any of the natural image databases we investigated, given fixed normalization factors. The
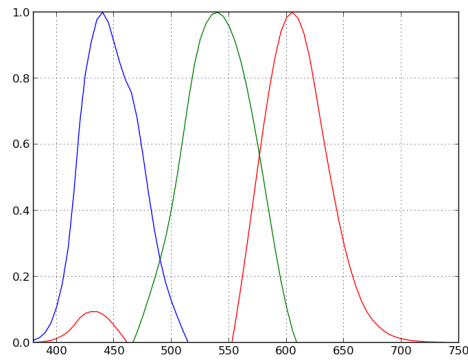
(a) sRGB
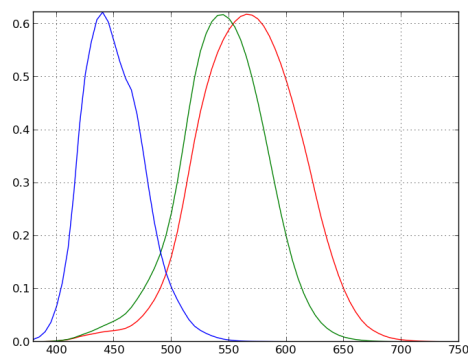
(b) LMS

Figure 5.3: *Photoreceptor sensitivities. The idealized model's sRGB cone sensitivities are more evenly distributed than the naive model's, so all cone correlations are approximately equal. All peaks are normalized to 1.0. Note that the cone sensitivities are a linear transform of the CIE 1931 standard observer cmf (figure 2.4), thresholded at 0, which causes a small peak in the short-wavelength region for the long-wavelength cone. If there were enough hyperspectral images of natural scenes, it would be possible to model any desired sensitivities without artefacts.*

De Paula (2007) model used a small set of images and manually chosen scaling factors on each channel in order to achieve a CR map in which all three cone types were represented. In the previous chapter, we saw that having equal overall channel activities can instead be achieved through homeostasis. We take equal overall excitations as the idealized position on this dimension.

In the second part of the results section, we will systematically investigate the effect of relative photoreceptor channel strengths. To explore the effect of imbalance, we form a dimension d_{CS} by starting from the idealized position (all channels having equal overall activity; $d_{CS} = 0.0$) and boost the first channel relative to the other two, making the activity ratio increasingly unbalanced compared to the idealized case. We only increase the *I* channel relative to the other two for simplicity in interpreting the results:

$$\begin{pmatrix} I_i' \\ J_i' \\ K_i' \end{pmatrix} = \begin{pmatrix} zI_i \\ \frac{3-z}{2}J_i \\ \frac{3-z}{2}K_i \end{pmatrix} \tag{5.3.6}$$

where $z = \frac{1}{3}d_{CS} + 1$, so that for $d_{CS} = 0$ the channels are just *IJK* from the previous section (i.e. idealized), and for $d_{CS} = 1$ the *I*'s scaling factor is 1.6 times that of *J* and *K*. (We will see later that this value gives more than enough difference between the channels to demonstrate the effect of relative photoreceptor channel strengths.)

Dimension d_{CS} lets us understand the effect of unbalanced channel activities, which could be caused by a cone type having a wider sensitivity (assuming the average incoming SPD is flat), or there being more of one cone type than another. A database with a non-flat average SPD could also cause unbalanced channel activities, but we have included this effect already in the input hue bias dimension.

5.3.4 Photoreceptors: correlation

We introduced the idea of photoreceptor sensitivity correlation earlier, using sRGB sensitivities instead of LMS for the idealized model. To form a dimension, which we will investigate in the second part of the results (section 5.5), we can start from the idealized position and subsequently increase correlation between one pair of photoreceptors by adding a fraction of one receptor type's activity to the other type's activity at each location. The amount of additional correlation between the pair of channels is

specified by d_C (which takes values from 0 to 1). Then for each location i:

$$\begin{pmatrix} I_i'' \\ J_i'' \\ K_i'' \end{pmatrix} = \begin{pmatrix} I_i' \\ (1-d_C)J_i' + d_C I_i' \\ K_i' \end{pmatrix} \tag{5.3.7}$$

At $d_C = 1.0$ the I'' and J'' channels are completely correlated (i.e. identical). At $d_C = 0.0$, the idealized position, the channels are just IJK' from the previous section.

Having calculated the possibly correlation-adjusted and possibly channel-ratio-adjusted photoreceptor excitations, we have now finished introducing the dimensions that apply to the photoreceptors. In the previous chapter, the excitations were transformed to photoreceptor responses including a model of homeostasis. However, in this chapter we want to investigate the effects of the dimensions, so we do not apply the homeostatic mechanism (i.e. no gains in equation 4.2.2). The idealized model has photoreceptor receptor activations that are already effectively normalized because of the full hue jitter (section 5.3.1), equal relative channel strengths (section 5.3.3), and equal correlations (above). Instead, we apply only σ (eqn 4.2.4) to have an approximately constant mean activation of the retina each iteration. This prevents any spurious effects caused only by changes in the overall amount of activity.

$$\eta_i = \sigma \begin{pmatrix} I_i'' \\ J_i'' \\ K_i'' \end{pmatrix} \tag{5.3.8}$$

5.3.5 RGC/LGN

To the RGC stage of the previous chapter's model we add the ability to:

- change the number and type of opponent channels (to investigate RGC pathway symmetry);
- set the LUM/CR ratio (to investigate the balance of luminance vs. CR).

The RGC processing pathway is outlined in figure 5.4.

5.3.6 RGC: pathway symmetry

As described in the previous chapter, the RGCs are modeled as Difference-of-Gaussians (DoG) connection fields (CFs), creating spatially and cone-opponent circular CFs.
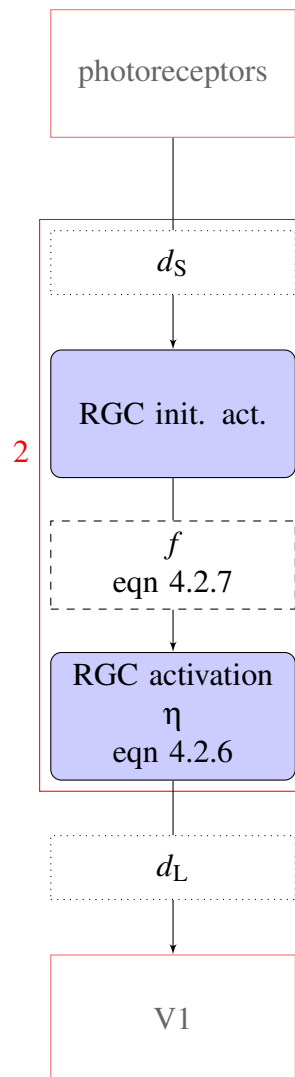
Figure 5.4: *RGC processing. The RGC pathway can either be symmetric, opposing all cone types equally, or it can be unsymmetric, as for the naive model—controlled by d_S. For training iterations, where V1's afferent weights to the RGC and LGN may change, the LUM:CR weight balance is enforced by d_L. Otherwise the processing is the same as for the naive model (figure 4.6).*

(a) Unsymmetric (naive)

(b) Symmetric (idealized)
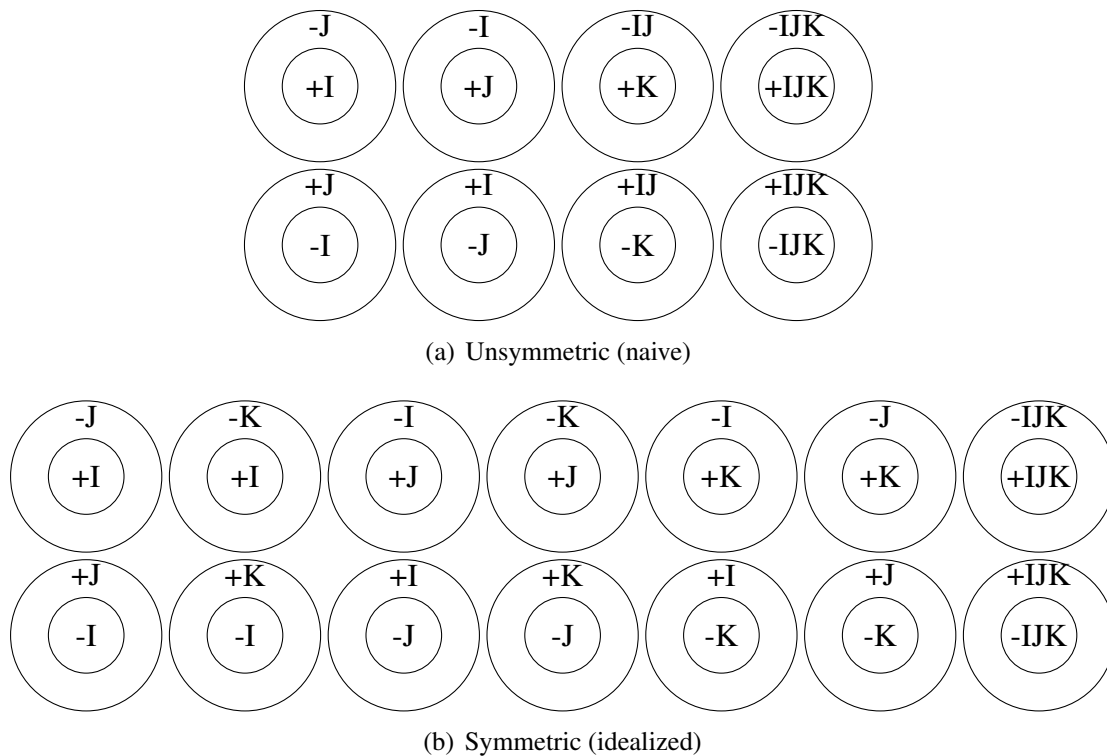
Figure 5.5: *RGC pathway. (a) Standard, currently accepted RGC classes (from the naive model). (b) Idealized, symmetric RGC pathway, with all combinations of pairs of photoreceptor types. In both cases, the luminance signal (right-hand column) is the sum of all cone types.*

As discussed in section 5.2, the effect of this pathway on hue representation in the model is not known, but in the naive model, not all opponencies are present. Therefore, we create the idealized situation of having every pair of possible photoreceptor combinations—the "symmetric" pathway, figure 5.5(b)—rather than the original LMS "unsymmetric" pathway (figure 5.5(a)). We also have a dimension of RGC symmetry, comprising symmetric (idealized model, $d_S = 0$) and non-symmetric (naive model, $d_S = 1$).

5.3.7   RGC: luminance/color ratio

Each V1 unit receives afferent input from several chromatic channels and an achromatic (luminance) channel, as previously shown in figure 5.5. Experimental results show macaque V1 CR-responsive cells are organized into patches, whereas OR maps are spatially contiguous. In the model, the balance of CR map and OR map can be controlled by setting the ratio of weights for the LUM vs. CR RGC pathways. Setting parameter $d_L = 0.0$ causes the luminance (LUM) channels to have zero weighting and

the CR channels to have full weighting, while $d_L = 0.5$ sets the weighting of LUM and CR channels to be equal, and $d_L = 1.0$ sets the weighting entirely to the LUM channels. The idealized model has $d_L = 0.4$ because this value was found to match existing data from macaque (discussed in section 5.5).

The afferent connections to V1 are normalized and scaled in two separate groups, one for the LUM connections and one for the CR connections. I.e., following equation 4.2.14 but adding a scaling factor, for afferent connections to V1 neuron q from neuron p in LUM connections (LUM On, LUM Off) we have:

$$w_{pq}(t) = d_L \frac{w_{pq}(t-1) + \alpha_L \eta_q \eta_p}{\sum_u (w_{uq}(t-1) + \alpha_L \eta_q \eta_u)} \qquad (5.3.9)$$

where u iterates over all LUM connections to q. For CR projections (I-J On, I-J Off, ...) we have:

$$w_{pq}(t) = (1 - d_L) \frac{w_{pq}(t-1) + \alpha_C \eta_q \eta_p}{\sum_u (w_{uq}(t-1) + \alpha_C \eta_q \eta_u)} \qquad (5.3.10)$$

where u iterates over all CR connections to q.

Additionally, for the results presented in this thesis, the learning rates were scaled for LUM vs. CR. For LUM connections, $\alpha_L = d_L \alpha_A$, while for CR connections, $\alpha_C = (1 - d_L)\alpha_A$. The learning rate scaling is for historical rather than functional reasons, and qualitatively similar results would be obtained if it were left fixed.

5.4   Results 1: Comparison with animal data

In the previous chapter, we compared the naive model to experimental data and found that while the naive model matched experimental results **ER1-ORmap** and **ER2-patches**, it did not match experimental results **ER3-range** and **ER4-peaks**. Therefore, the first goal of this chapter is to determine whether a model of this kind can develop a realistic representation of CR at all. To do this, we introduced the idealized model in the previous section. We will now evaluate the idealized model against experimental results **ER1-ORmap**, **ER2-patches**, **ER3-range**, and **ER4-peaks**. We find that the model can indeed match those experimental results. To our knowledge, this is the first time all these results have been matched by a developmental model. In a subsequent section, we will explore the idealized model's dimensions in order to understand the effect of the CR biases, and therefore what a rescue of each dimension would need to achieve. Doing so lays the groundwork for implementation of rescues in the future.

Earlier in this chapter, we identified aspects of the naive model that could bias the hue representation and control the spatial representation of CR vs. OR, and then we created a model with an appropriate balance of CR vs. LUM and without these sources of bias to match experimental data—the idealized model. We now analyze this new model using the same analyses as in chapter 4.

5.4.1   Joint representation of OR and CR

The idealized model's representation of OR is very similar to the naive model's. Rows two (OR sel), three (OR pref), and four (OR pref & sel) of figure 5.6 show a comparison of the idealized model to the naive model and experimental results. We can see that the model develops an OR preference and selectivity map consistent with the variety of experimental data. Hence the model continues to match **ER1-ORmap**.

Second, the idealized model shows similar CR patches to the naive model in response to RG-BW gratings (figure 5.6, row 1). As for the naive model, differences in CR patch size and spacing between the model and experiment are no greater than differences between experiments. Also as for the naive model, the CR patches tend to occur in areas of lower OR selectivity, as in experimental data. Therefore the model also continues to match **ER2-patches**.

5.4.2   Representation of hue

The main question is whether the model can develop a realistic representation of hue, i.e. one that allows hue to be coded by spatial position or preference. Figure 5.7 shows results from the idealized model compared to experimental data and the naive model. We can see that the idealized model does indeed develop a realistic representation of hue, unlike the naive model. Rows 1 (CR peaks) and 2 (CR contours) show that patches of high CR selectivity represent many or all hues (first column), rather than a limited range (naive model; second column). Hence the idealized model matches **ER3-range**, so we have solved problem **P1-range**.

We can also see in row 1 (CR peaks) that the peaks are spatially well separated. Row 3 (CR peak distances) confirms this, and shows that the distance on V1 corresponds to the distance in PCS (including the circular nature of PCS: the peak responses to red and violet are closer to each other than either is to green). Hence the idealized model matches **ER4-PCS**, so we have also solved problem **P2-PCS** with previous modeling work.

Finally, examining the last three rows of figure 5.7—the model's CR preference and selectivity maps—allows us to understand better the model's CR organization. All hue preferences are equally represented, and hue is organized in a map that is (a) less dominant than the OR map (hence giving the patchy CR responses) and (b) larger scale than the OR map (the Fourier transform—not shown—has a larger ring diameter). This addresses **P3-pref**: all hues are now represented equally. The V1 hue preference histogram matches the input histogram.

Having seen these results, we can therefore conclude that it *is* possible for this kind of developmental model to achieve a realistic representation of CR and code hue by place or preference. The previously listed CR-biasing dimensions (input hue bias, photoreceptor correlation, relative photoreceptor channel strengths, RGC pathway symmetry, and luminance vs. CR) therefore seem important, and likely hold the clue to understanding problems **P1-range**, **P2-PCS**, and **P3-pref**. However, to demonstrate that each one of these is critical, and to understand what role each plays, we will now consider each of the above individually, as independent dimensions. What separates the naive model from the idealized model on each dimension, and what would it take to rescue the naive model to achieve realistic results for that dimension? And, as mentioned before, another reason to investigate these changes as dimensions is that different species and/or individuals could lie in different positions, and the results could be used to predict the results of specific rearing conditions, physical abnormalities, or experimental treatments.

Figure 5.6 *(following page)*: *Comparison of idealized model, naive model, and experimental data for CR vs OR organization. Experimental data was introduced in same format in figure 2.20: each 1.5 inch square is 3 mm x 3 mm of V1, and NA indicates no data available. From the first column, we can see that the idealized model forms a spatially contiguous OR map, and that the responses to RG-BW remain patchy and tend to occur in regions of low OR selectivity, as for experimental data. The idealized model's OR vs. CR representation is as good as the naive model's; these manipulations primarily affect details of the CR representation (which is why these aspects have been ignored in previous luminance-only models).*
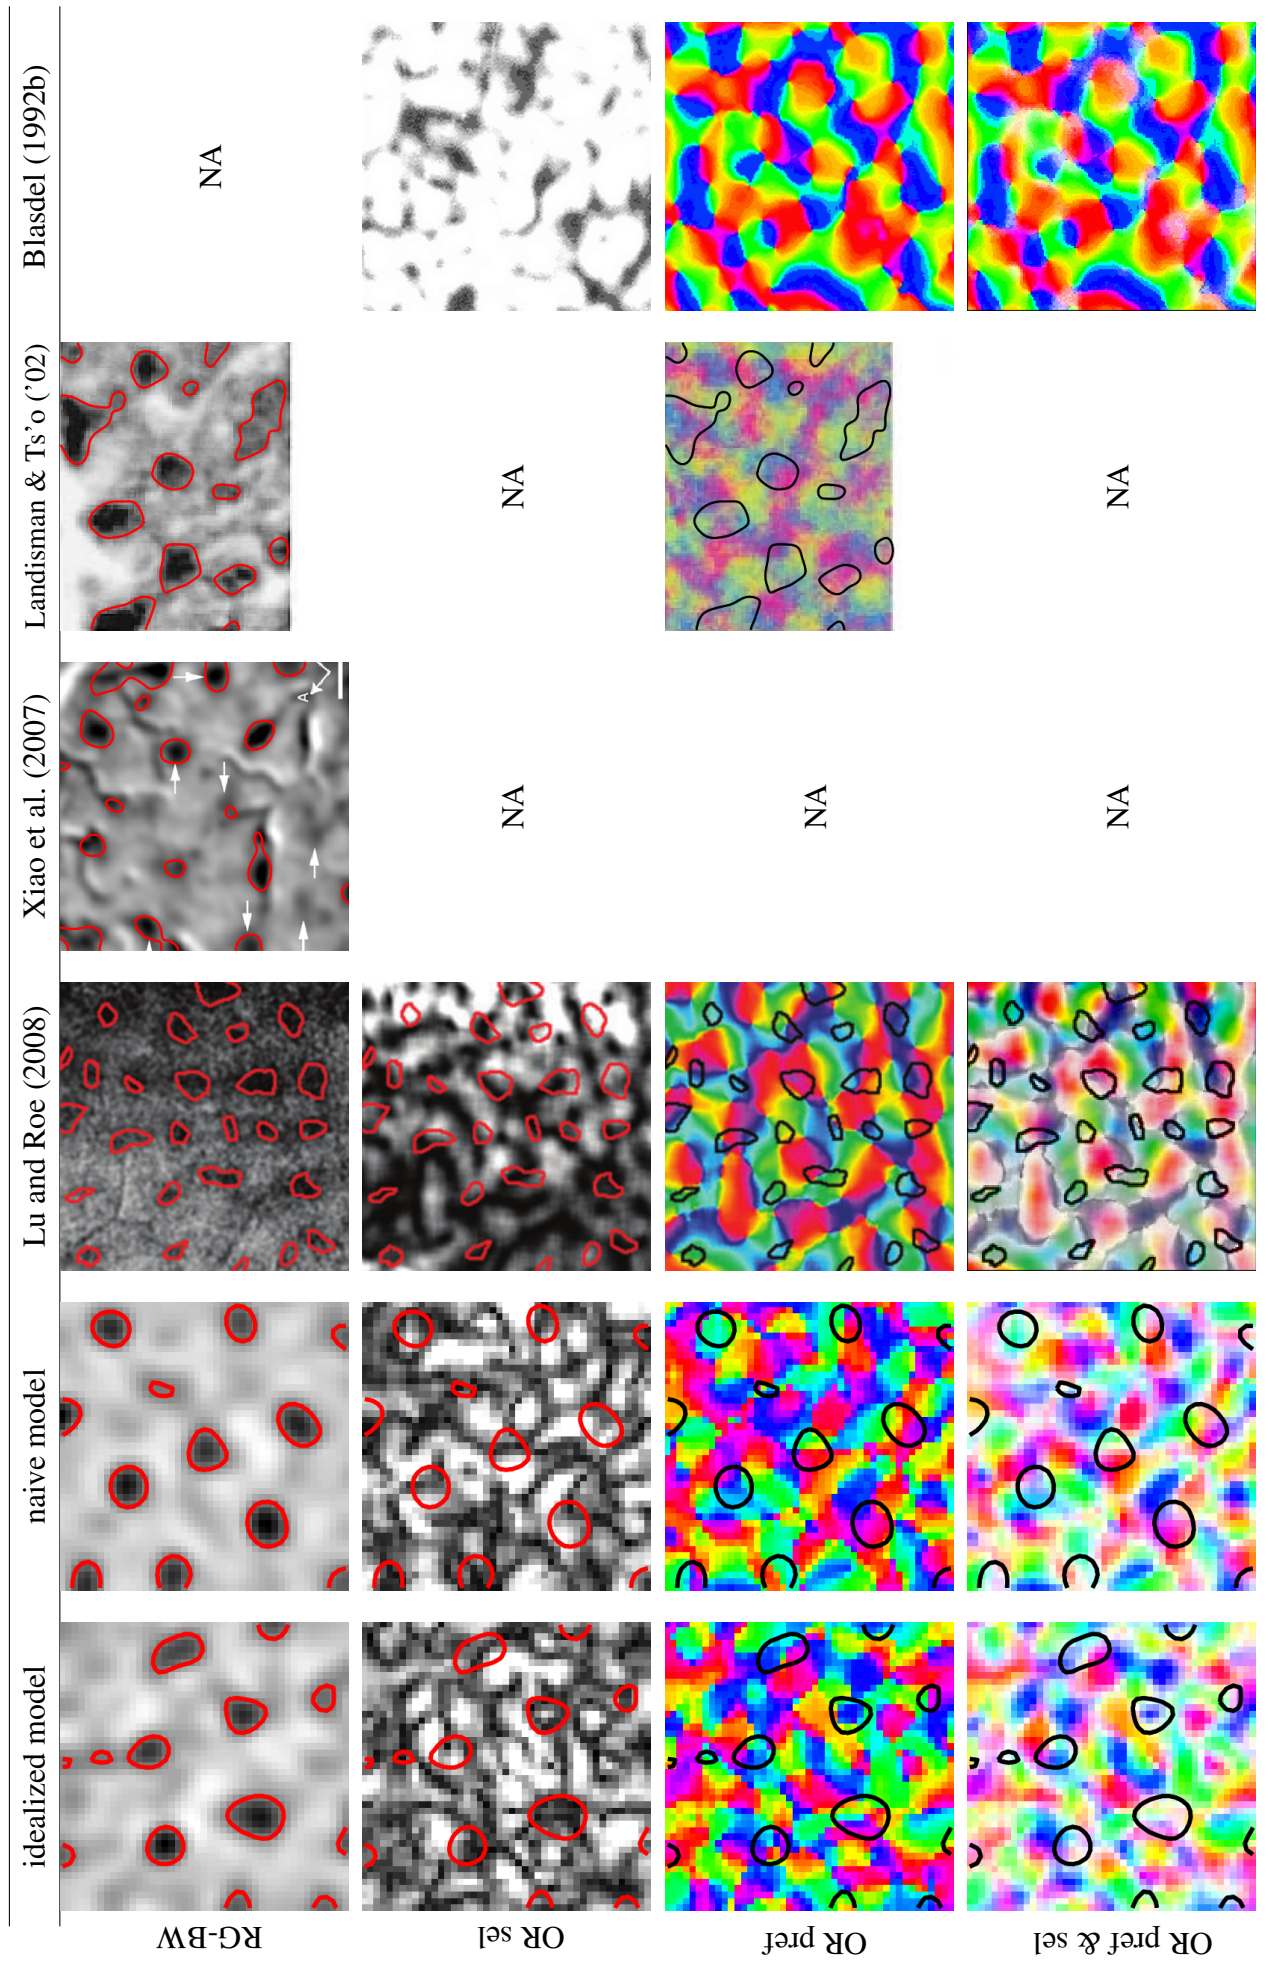
	ER1-ORmap	**ER2-patches**	**ER3-range**	**ER4-PCS**
De Paula (2007)	✓	✓	✗	✗
Naive model	✓	✓	✗	✗
Idealized model	✓	✓	✓	✓

Table 5.1: *Models vs. experimental claims. Among existing models of the development of OR and CR (reviewed in chapter 2), only the three shown above are able to account for the simultaneous development of OR and CR into organizations that match data from macaque V1. The naive model extended the work of De Paula (2007) to improve the model's development process (including obtaining higher quality OR maps), reduce the number of free parameters, and compare more closely to experimental data. However, the naive model was unable to account for the organization of hue. The idealized model is able to develop a realistic organization of hue in addition to OR and OR vs. CR.*

5.5   Results 2: Dimensions

In the previous section, we showed that a model with dimensions set to idealized positions (the idealized model) could develop a representation of CR matching experimental data. That achieves the first goal of this chapter: showing a model of this kind can represent hue in a way that allows coding by preference or spatial location at all. However, we still need to show that each dimension critically affects the model's de-

Figure 5.7 *(following page)*: *Comparison of idealized model, naive model, and experimental data for hue organization. Experimental data was introduced in this format in figure 2.23: each 1.5 inch square is 3 mm x 3 mm, and NA indicates no available data. First three rows: regions of high CR selectivity in the idealized model contain a hue organization matching experimental data. Patches respond to a wide range of hues, and within a patch the responses to different hues are spatially well separated. In particular, row 3 shows that the distances between peaks on the cortex correspond to distances in PCS (both model and experiment represent red at 0 on the x-axis; the model's scale should be multiplied by 180 to convert to degrees). This contrasts with the naive model, in which patches respond to a limited range of hues, and peaks are not well separated. Last three rows: CR selectivity and preference maps—data from the model that is not yet available from animal experiments. The idealized model forms a CR map containing units that prefer all the different hues. When tested with RG-BW gratings, the response is patchy (red overlay on CR sel plot), with units preferring red through green responding the most. The naive model, on the other hand, as was originally shown in figure 4.11, develops patches which represent only a limited range of hues, and the peak hues are not well separated, making its hue representation unsuitable for building PCS. The idealized model is the first developmental model to demonstrate this realistic organization for perceptual color.*
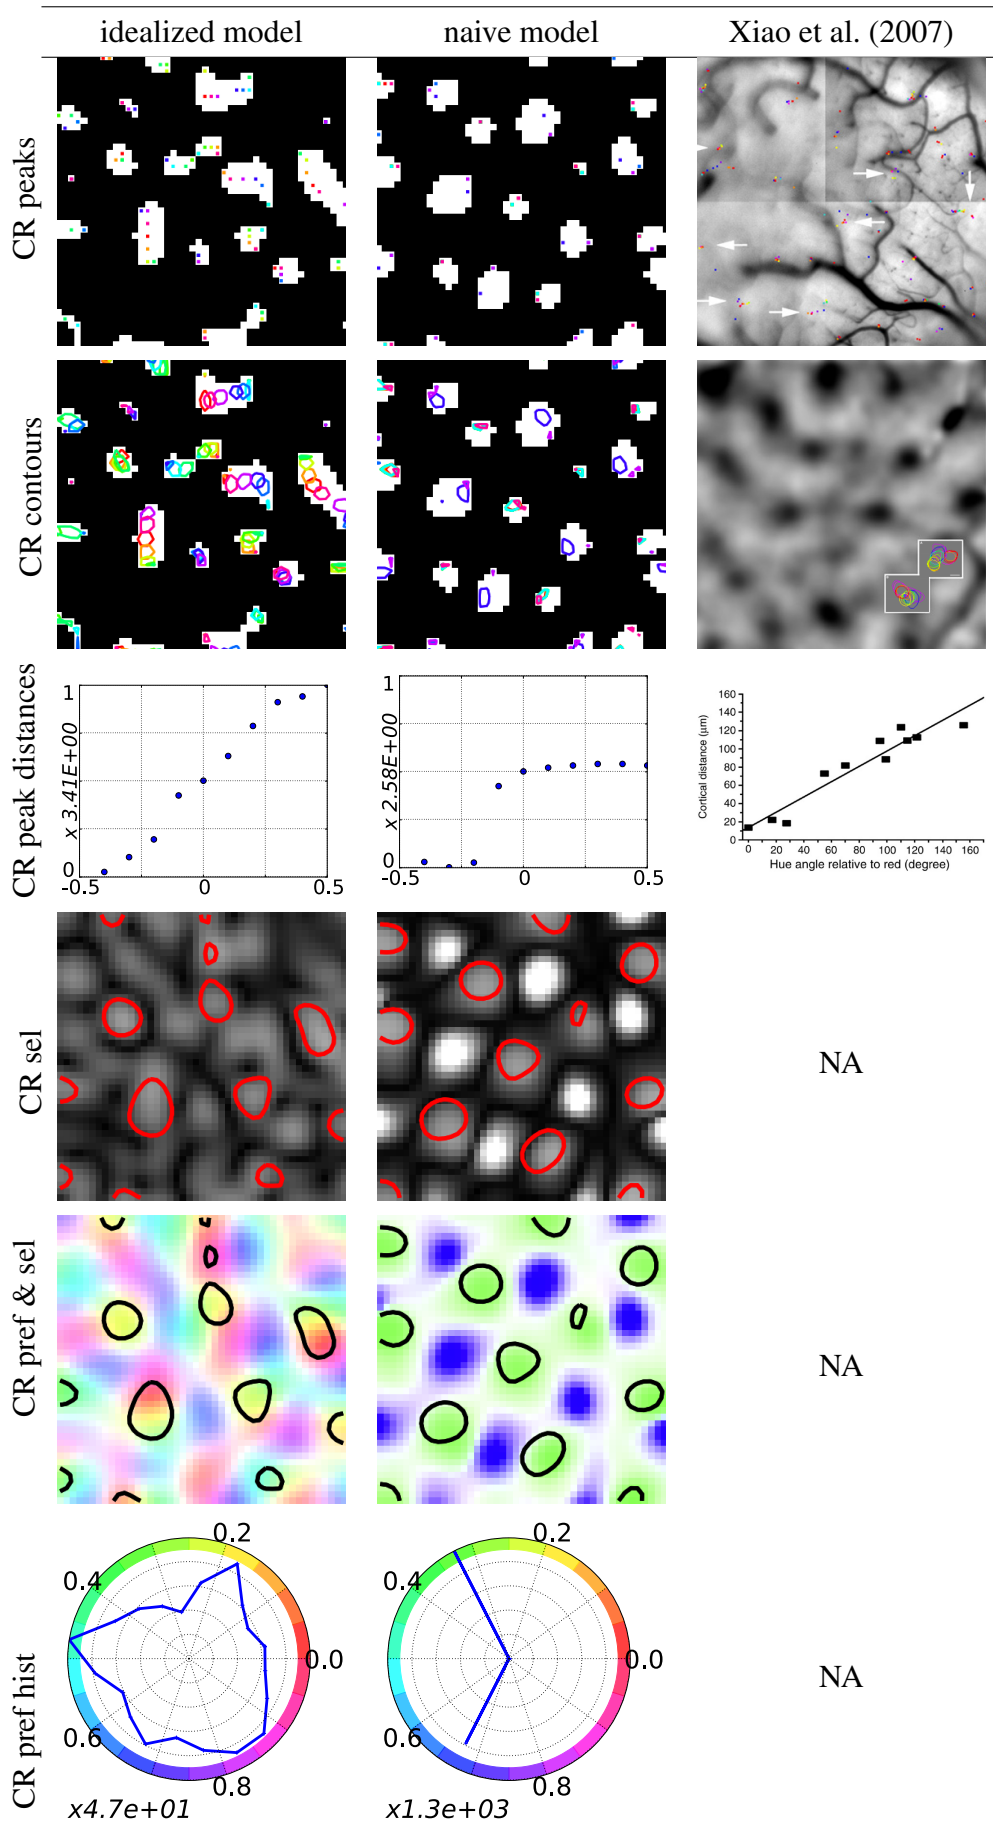
Dimension	Parameter	Idealized
1: Input hue bias	d_{HB}	0.0
2: Relative photoreceptor channel strengths	d_{CS}	0.0
3: Photoreceptor correlation	d_C	0.0
4: RGC pathway symmetry	d_S	0.0
5: Luminance vs. CR	d_L	0.6

Table 5.2: *Idealized model's dimensions. All dimensions have range 0.0 to 1.0.*

velopment. Therefore, in this section, we will show the results of manipulating each dimension independently. To do this, we will start from the idealized model, and re-introduce one dimension at a time (e.g. photoreceptor correlation) to varying degrees in order to determine its effect. This will additionally allow us to better understand the dimension (i.e. the problem the brain is apparently overcoming in developing PCS), and to inform what a rescue for each dimension would need to achieve.

In the following sections we will go through each of the dimensions, summarized in table 5.2. First, however, we will show the inherent variability in the simulations as a baseline against which we can compare the magnitude of the dimensions' effects.

5.5.1   Inherent variability

To see the impact of other dimensions, first we show the inherent variability in the simulations as a baseline. To do this, the hue jitter's seed is randomized, resulting in a different order of presentation of the hues (but ultimately the same uniform incoming distribution). The results in figure 5.8 show that while details of the OR and CR organizations vary, the overall statistics are similar. Development always leads to a hue preference map representing all hues, with the CR peaks plot showing most or all hue peaks per patch, and with cortical distance corresponding to distance in PCS. Figure 5.8 also introduces a number of plot types that will be used repeatedly in later figures.

An interesting observation is that there is some variety in e.g. the distribution of hue peaks in the patches (row CR pref hist). Although in all cases the underlying hue preference maps are well organized, and are overall quite similar, the CR patches (taken at a particular threshold of the CR selectivity) do not always have the same level of uniformity. This effect could be one source of differences between experimental results, with some studies seemingly finding patches dominated by one or a few preferences while others show many hues per patch. We will return to this issue later in the discus-

sion and following chapter.

Figure 5.8 *(following page)*: *Varying hue normalization's input seed. Each column represents a different simulation run, changing only the seed of the uniform random generator U, resulting in a different stream of hues. The seed is shown at the top of each column.*

This figure allows us to see the inherent variance: a different ordering of input hues leads to different details in the maps, but similar overall statistics. The left-hand column is the default random seed, as used in other simulations in this chapter. The rows, introducing chart types used in subsequent figures, are as follows.

Top row, **in hue**: *cumulative incoming hue distribution. The polar axis has HSV hue bins; hue counts over all input patch presentations throughout the whole duration of the simulation are shown on the radial axis. The radial axis limit is displayed at the bottom left; each tick is 0.2. Here, all incoming distributions are uniform.*

PR abs: *bar chart of cumulative total excitation of each photoreceptor type over the whole duration of the simulation. The y-axis limit is shown on the left. In this case, all classes are equally activated over the course of the simulation.*

lgn acts: *bar chart of cumulative total activation of each RGC type over the whole duration of the simulation, weighted by d_L. Key: each digit refers to a pair of OFF center ON surround, ON center OFF surround opponent channels: 0=K/J, 1=K/I, 2=J/K, 3=J/I, 4=LUM, 5=I/K, 6=I/J. Here the LUM channels have a higher strength than the cone-opponent channels, to get a realistic balance between the OR and CR maps.*

OR sel: *OR selectivity, as introduced earlier. The exact pattern of orientation selectivity varies with the input seed, but the overall properties are similar. Orientation preference is not shown because it is similarly unaffected.*

CR sel: *CR selectivity, as introduced earlier. Again, the exact pattern varies with input seed, but is overall similar.*

CR pref: *CR preference and selectivity, as introduced earlier.*

CR pref hist: *The distribution of preferences is reasonably uniform. Note that maps from simulations with a lower d_L (i.e. more weighting to chromatic channels) have more uniformly distributed preferences (shown in a later figure).*

CR peaks: *Location of peak response to each hue within patches, as introduced earlier.*

peak dist's: *Average distance of each peak from red (0.0) within patches, introduced earlier.*

CR contours: *Fixed fraction of the peak response to each hue within patches, as introduced earlier.*

CR responses: *Radial axis shows total V1 response to a full-field, spatially uniform pattern of the hue indicated on the polar axis. In this case, V1 responds a similar amount to every hue.*
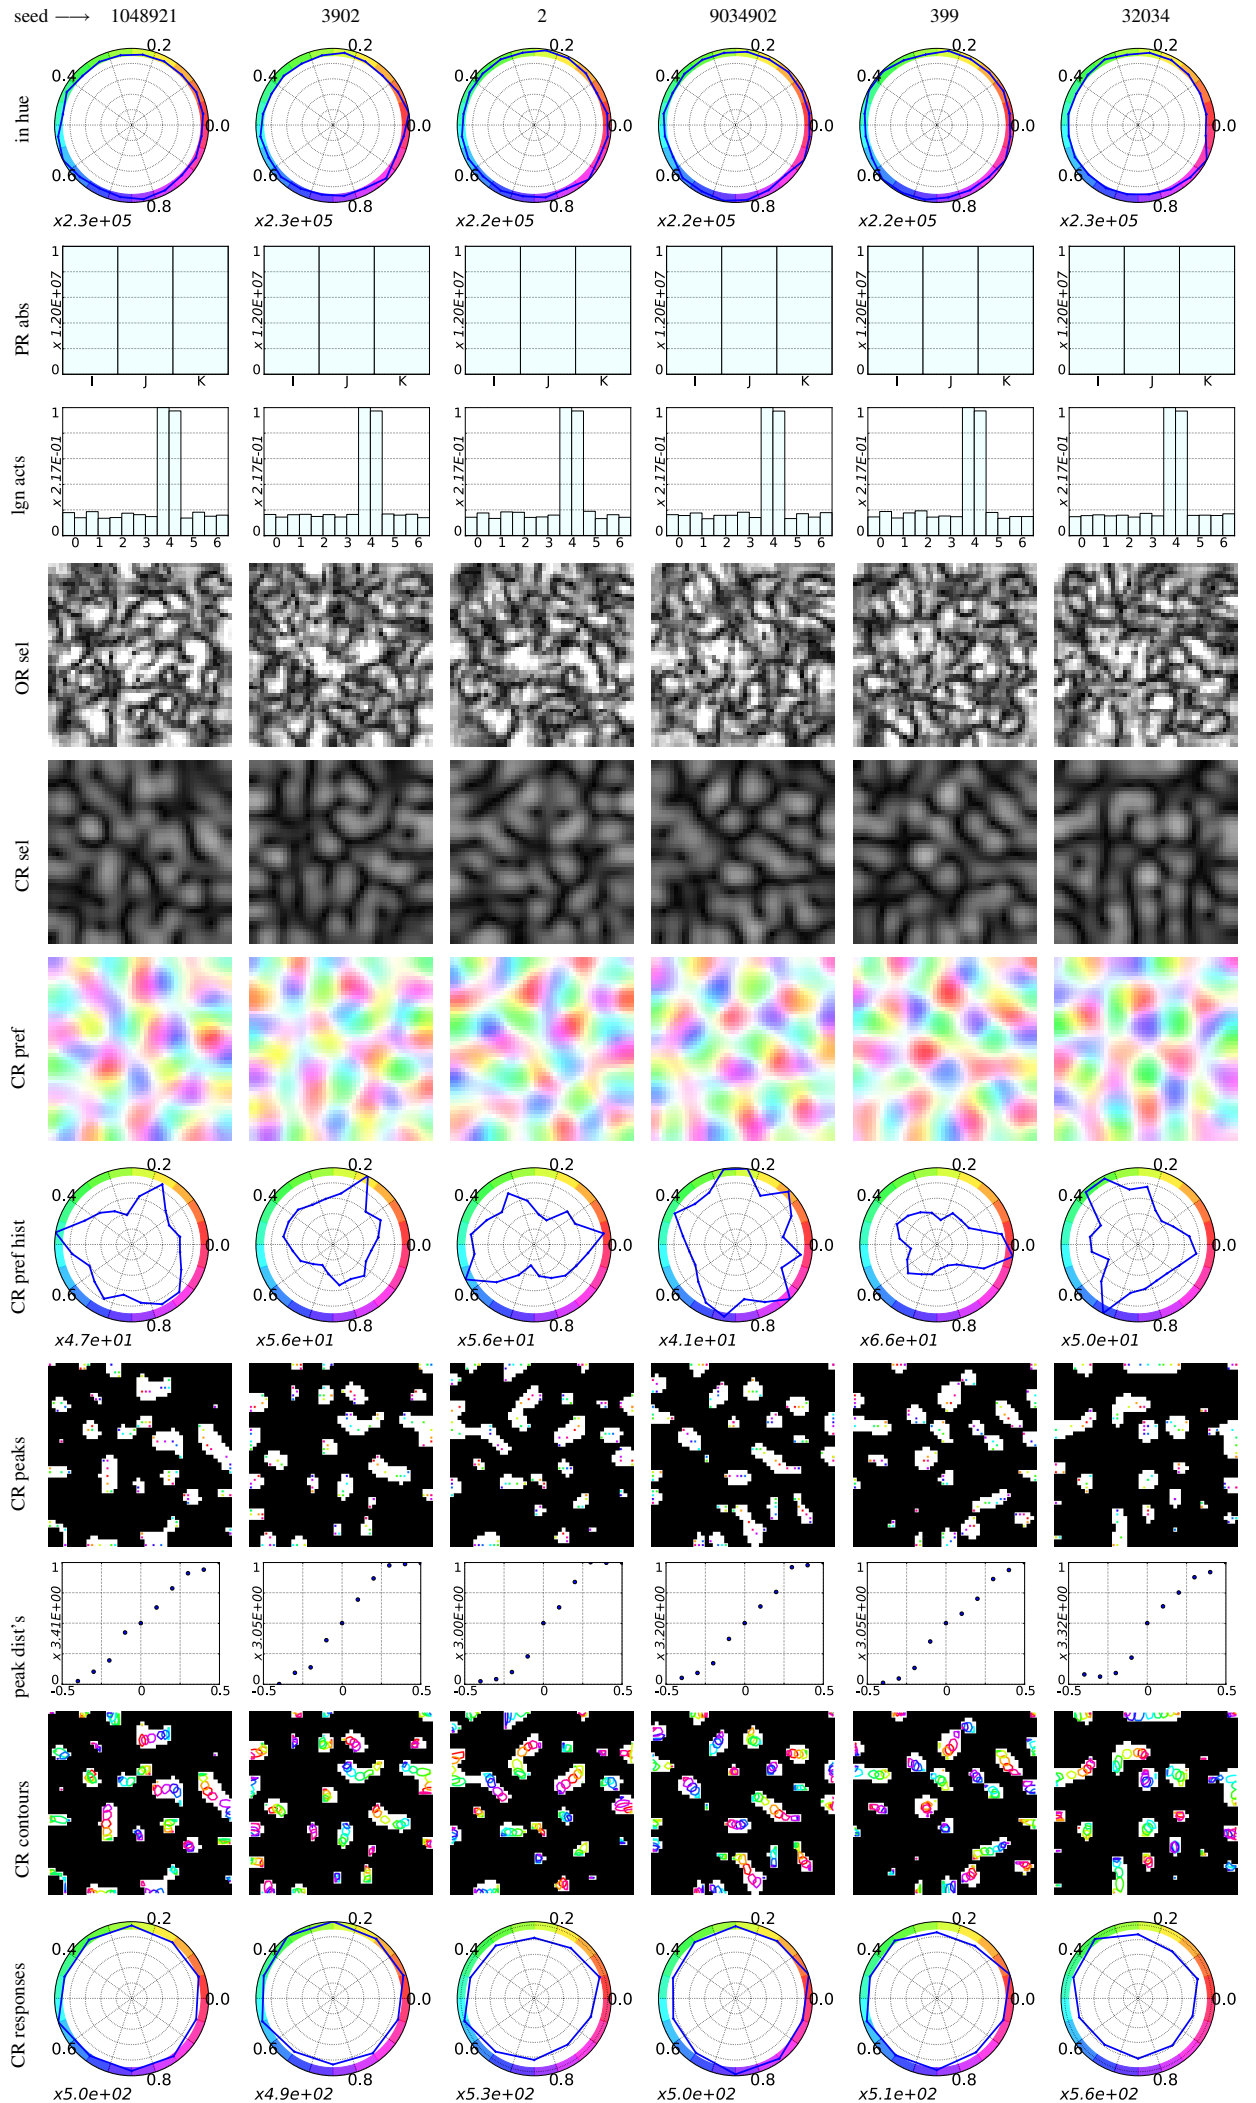
5.5.2   Dimension 1: Input hue bias

Chapter 3 showed that the distribution of perceptual hues in currently available calibrated color natural image databases is a highly non-uniform sampling of PCS. For the idealized model, we showed that jittering the hues in an image database to give a uniform sampling of PCS resulted in V1 learning all hues (solving **P3-pref**). To understand this better, we can investigate the results of gradually increasing the amount of bias, from the idealized, unbiased case ($d_{HB} = 0$) through to the full bias/original distribution ($d_{HB} = 1.0$).

Figure 5.9 shows the effect of increasing the input hue bias. While OR organization is unaffected, CR organizes very differently; the map becomes dominated by neurons preferring hues around the input bias. In fact, V1 amplifies the incoming bias. CR patches become larger, but despite this, many hues are not spatially well separated. If some hues (i.e. ratios of IJK photoreceptor activations) are not present in the training data, or are present only infrequently, V1 does not represent them. In our model, without any form of adaptation, cone activation ratios correspond one-to-one with hues.

Additionally, increasing the hue bias in the incoming data has another effect. In this example, the hue bias results in less energy at short wavelengths in the database, and therefore the K and J cones are activated less than the I cones. We already know how to solve this aspect of the problem from the previous chapter (i.e. some form of adaptation that effectively normalizes the photoreceptor activations), and we will consider it as part of a separate dimension (relative photoreceptor channel strengths).

Hence, this dimension contributes to all of the *PCS representation problems* (problems **P1-range**, **P2-PCS**, and **P3-pref**). Somehow, the early visual system is able to develop an apparently uniform representation of PCS, despite the environment being a highly non-uniform sampling of PCS. This non-uniform sampling may be occurring over very long timescales (i.e. a lifetime spent in yellow, arid desert), or over medium timescales (e.g. changing season, from lush green to arid yellow), and additionally even over short timescales (e.g. changing daylight over the day, or from weather changes, or even just from moving between locations). That is, although we may think of this dimension as being mainly about the reflectances present in the environment, the visual system actually receives radiances, and these radiances will vary dramatically in different situations. In this model, the results will be a highly biased and incomplete representation of hue (row **CR pref hist** and **CR responses**).

Figure 5.9 *(following page)*: *Dimension 1: input hue bias. Top row: the amount of incoming hue bias is increased from $d_{\mathrm{HB}} = 0.0$ (idealized, uniform distribution; left-hand column) to $d_{\mathrm{HB}} = 1.0$ (original, biased distribution; right-hand column). The hue preference map (CR pref) and histogram (CR pref hist) show that the model reflects and amplifies the incoming hue bias. The bias in the model appears to be the consequence of two factors. The first is that photoreceptor excitations fall for J ("medium") and K ("short") cones (shown in the PR abs row), as there is now relatively more energy at longer wavelengths. This causes fewer V1 units to respond at all to short-wavelength inputs, as shown in the last row (CR responses). However, even if the photoreceptors are normalized by long-term adaptation (as introduced for the naive model in the previous chapter), the hue bias has an additional effect: some ratios of photoreceptor activations occur more often than others, so the Hebbian-learning V1 represents these more strongly (i.e. with more units) than the patterns that occur less often. This is something we will return to in the next chapter.*
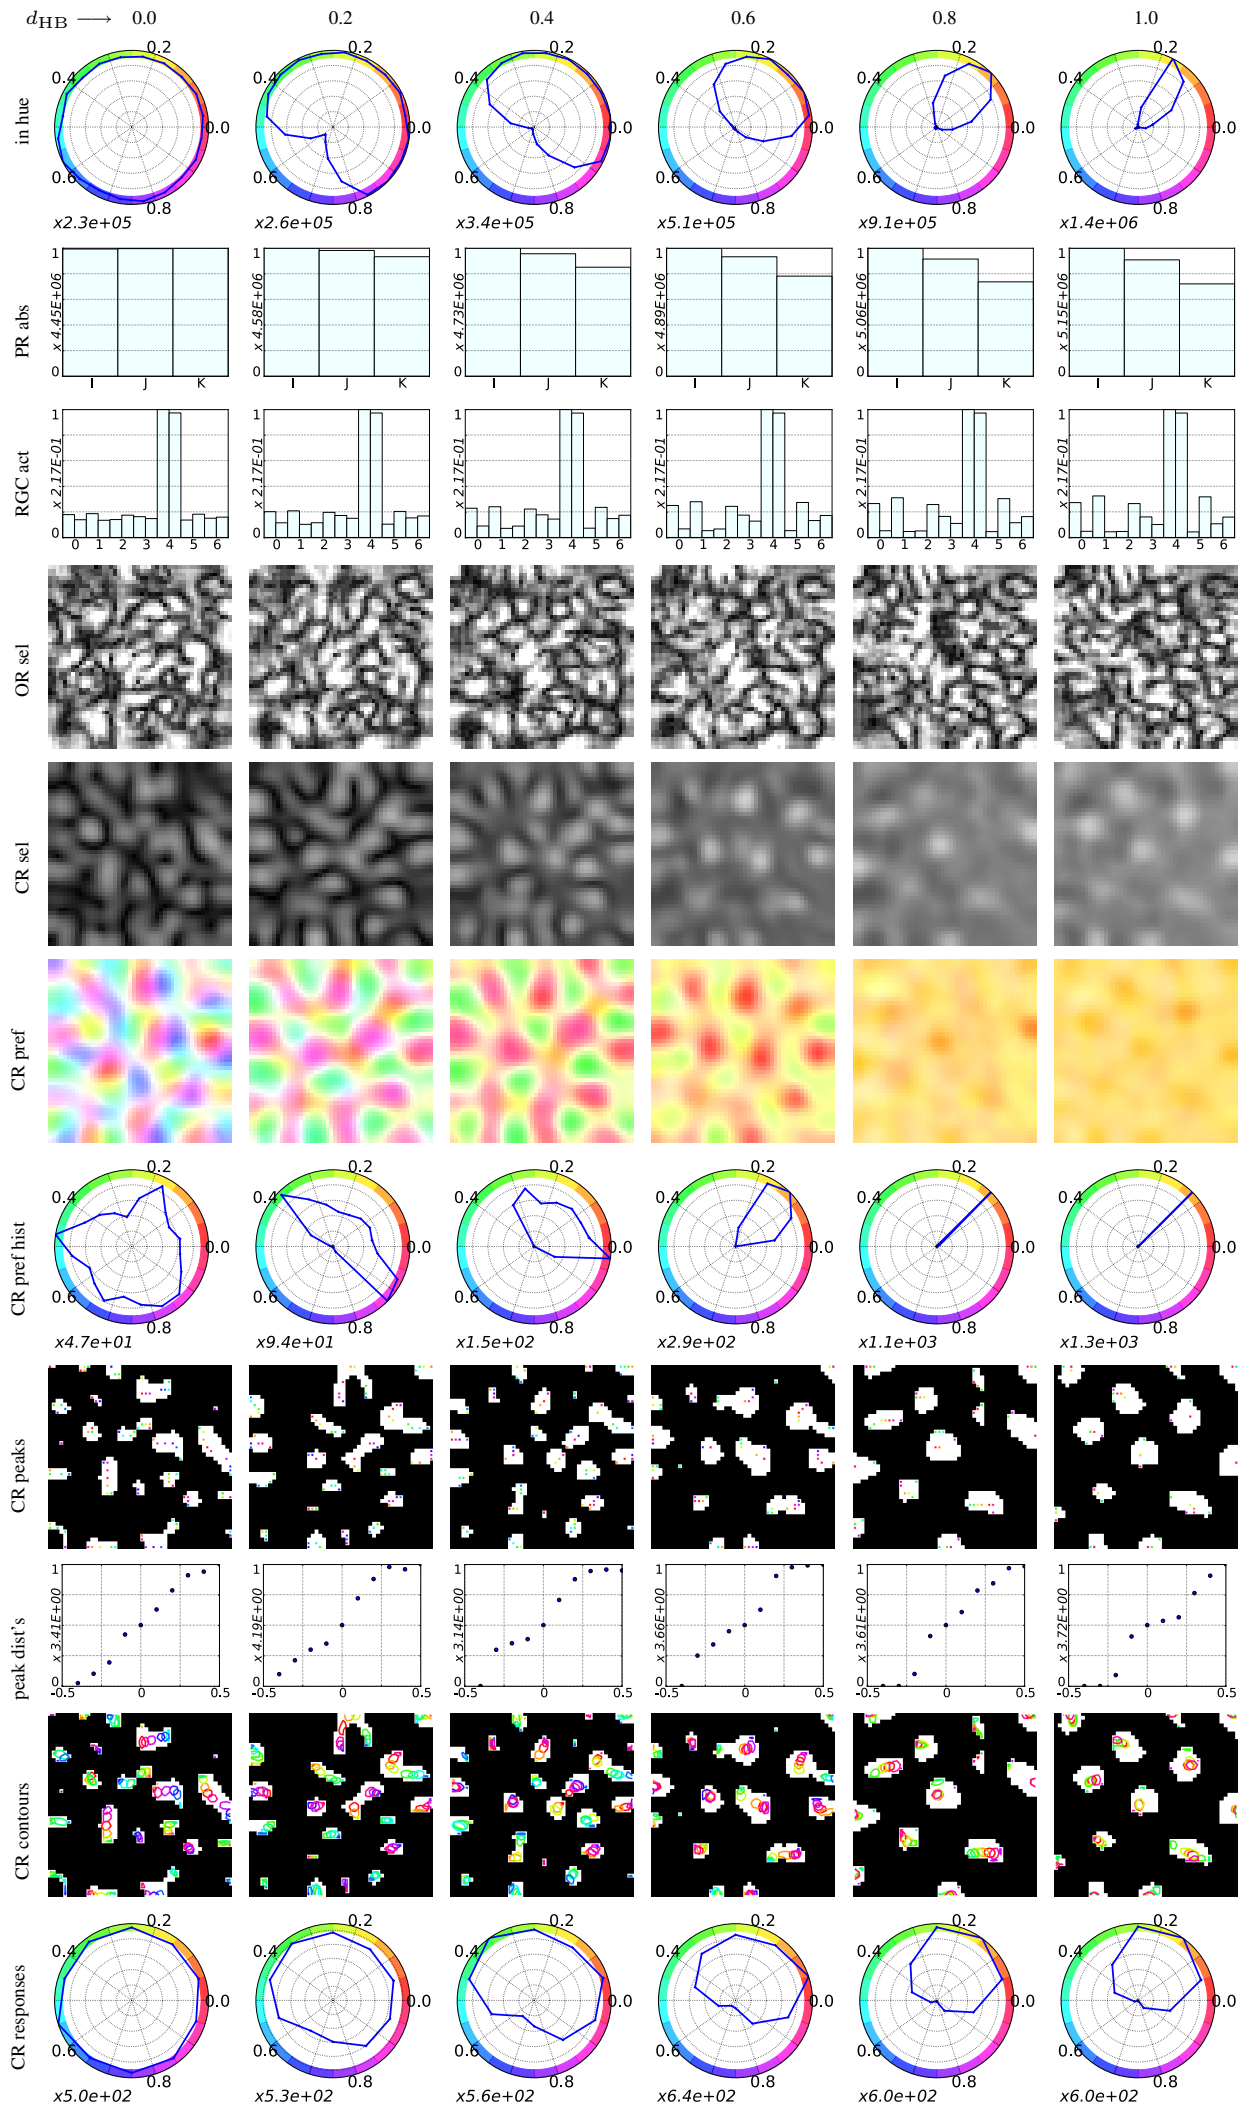
5.5.3 Dimension 2: Relative photoreceptor channel strengths

Here we consider uneven activity of the different photoreceptor classes, which could be caused by the distribution of wavelengths in natural scenes relative to cone sensitivities (as mentioned above), but also by variable ratios of cone types, and variable spectral sensitivities (e.g. the lens absorbing increasingly more short-wavelength light with age). Chapter 3 showed how the total excitations of the L, M, and S photoreceptors are not equal for any given image database converted to L, M, and S cone excitations using a fixed normalization factor. In chapter 4, we already addressed ("rescued") this problem with adaptive homeostasis at the photoreceptor level. However, to see the effect of this dimension and what the rescue achieves, here we simulate the effect of this dimension by boosting one of the relative photoreceptor channel strengths relative to the other two.

The position along this dimension is determined by parameter d_{CS}; the top row of figure 5.10 shows that when d_{CS} is increased from 0 through to 1, the I channel is boosted relative to the other two. The preference map (OR pref & sel) and histogram (CR pref hist) show preferences shrink until eventually the model comes to be dominated by the I channel, i.e. red preferences. Additionally, the CR responses show that V1 now responds to a smaller and smaller range of hues until eventually it only responds at all to a limited range of hues around red.

We can conclude that without any activity-based normalization, if one cone type is more active than the others, the competitive Hebbian learning V1 will become dominated by that cone type. Therefore, if a database contains more energy at L cone wavelengths than at S cone wavelengths, or if we were to simulate fewer S cones than L cones, or for any of the reasons mentioned earlier that would cause one cone type to be overall less active than another, then without some form of compensation, V1 will be dominated by the more active cone type. Therefore, a rescue of this dimension needs to equalize the activities between the channels—which adaptive homeostasis at the photoreceptor level does, for equal cone ratios. This dimension clearly also contributes to problems **P1-range**, **P2-PCS**, and **P3-pref**, since the dominance of one channel prevents some hues from being represented in V1 at all.

Figure 5.10 *(following page)*: *Dimension 2: relative photoreceptor channel strengths. The I channel is boosted relative to the other two channels, from $d_{CS} = 0.0$ (no boost; idealized model) in the left-hand column to $d_{CS} = 1.0$ in the right-hand column. This dimension simulates uneven activity of the different photoreceptor classes, caused by the distribution of wavelengths in natural scenes, variable cone sensitivities, and relative numbers of cone types. When the I channel is boosted relative to the other two, the CR preference map becomes dominated by hues that depend on I channel input. Boosting the red channel relative to the other two results first in preferences shrinking around the bias, then eventually in all neurons preferring red and responding to only a limited range of colors around red. This dimension confirms our findings from the previous chapter with the naive model. Also interesting to note is that the CR selectivity rises (and patch sizes increase), and OR selectivity falls, as the I channel is boosted. The change in balance of CR and OR in V1 happens because the relative strength of the incoming CR signal (vs. luminance) increases, since there is increasingly often a positive difference between the I channel and the other channels. The RGC act bar charts show this difference—an increase in activation of K ON center/I OFF surround (1), J ON center/I OFF surround (3), I ON center/K OFF surround (5), and I ON center/J OFF surround (6) channels.*
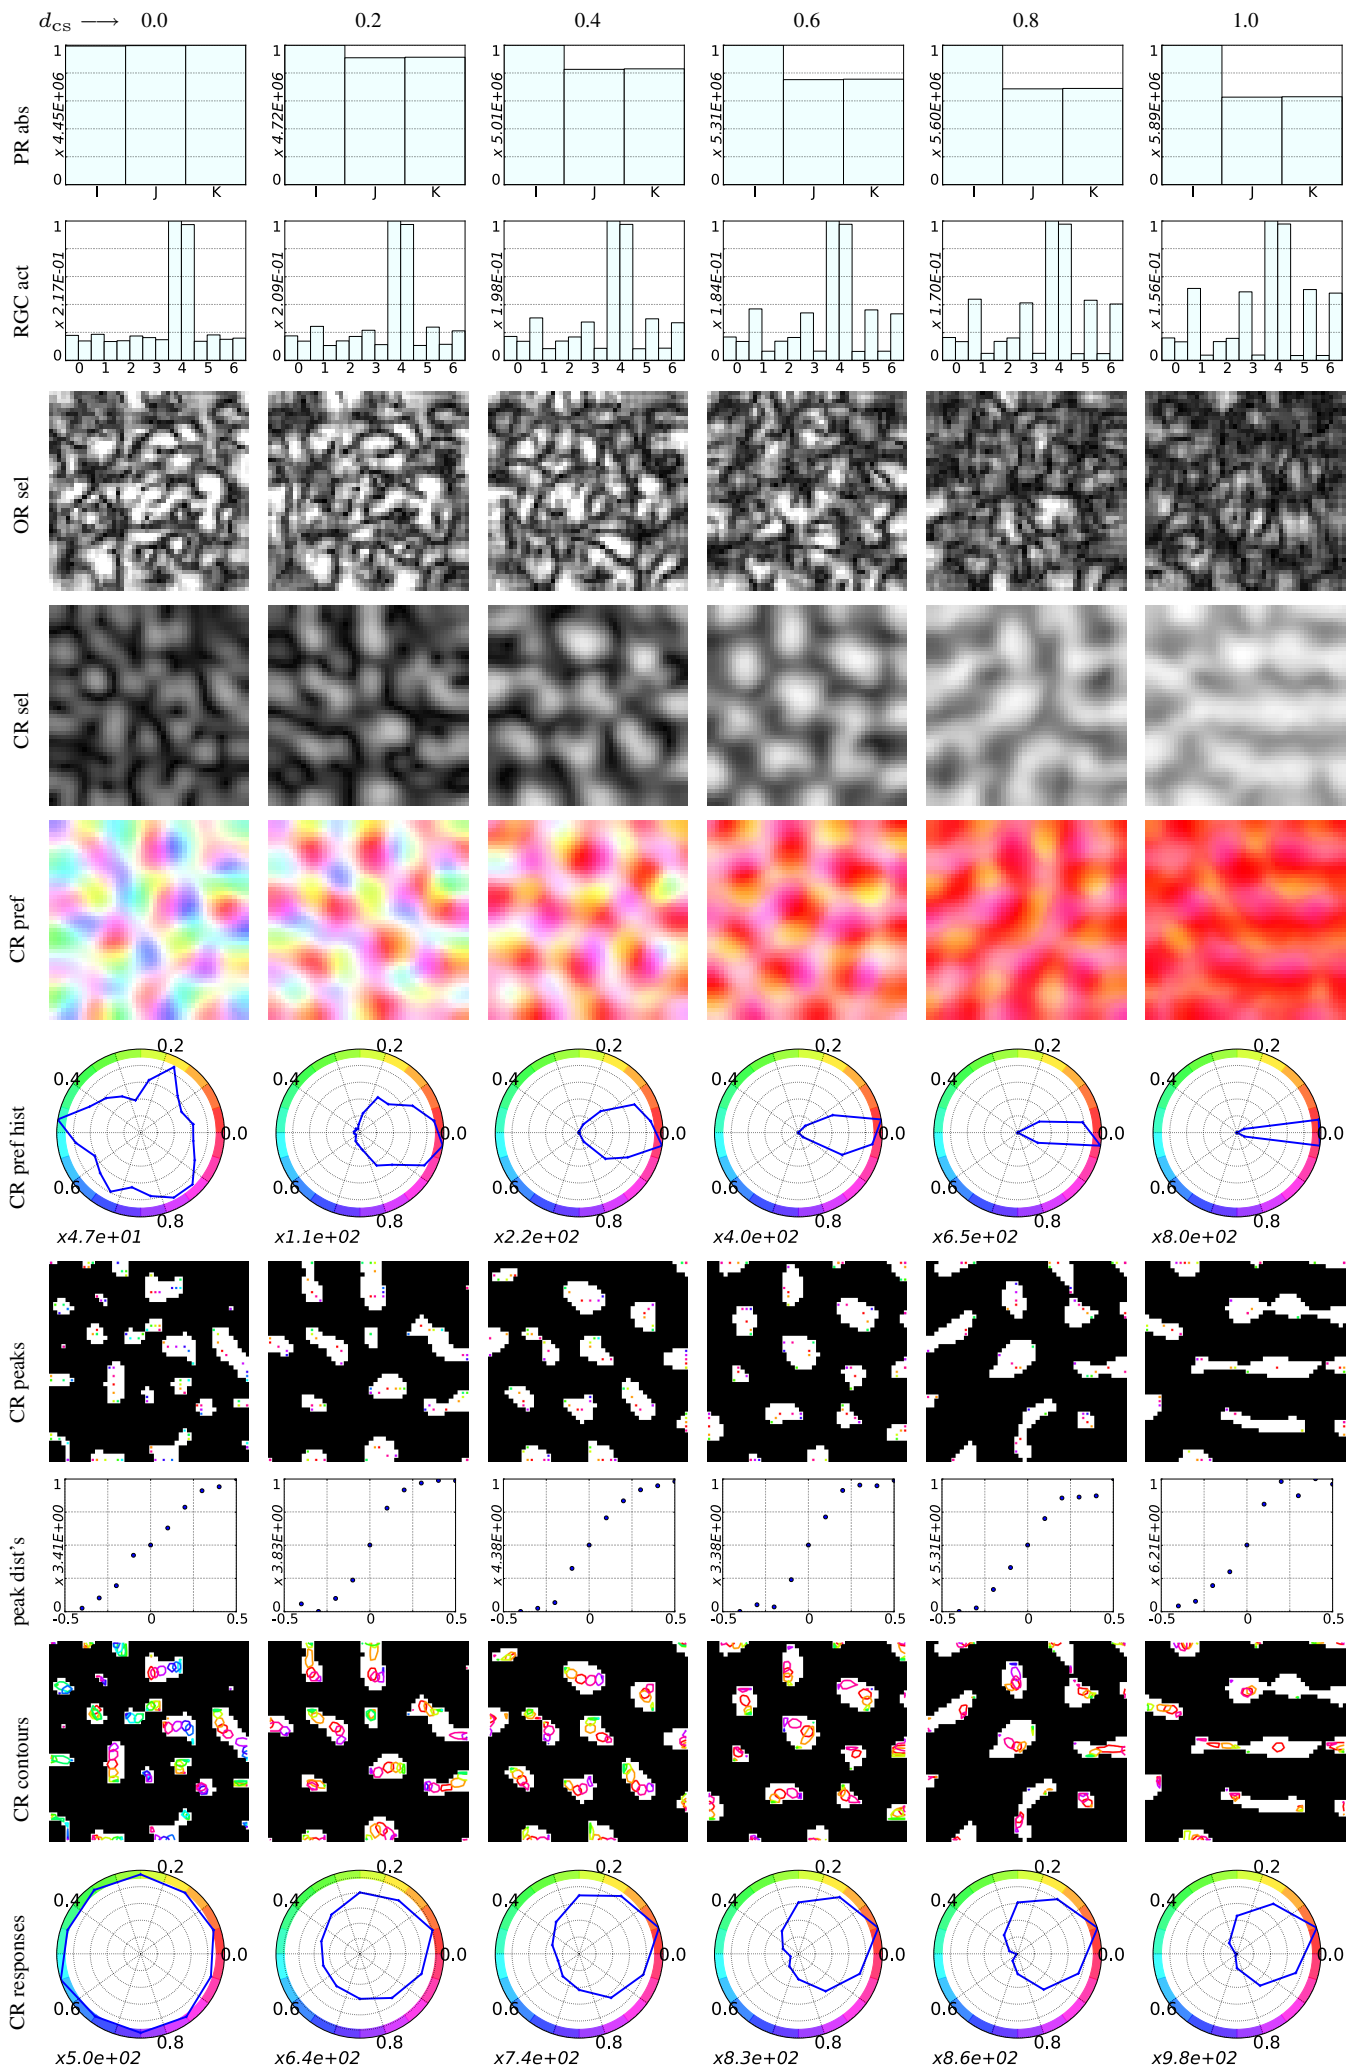
5.5.4   Dimension 3: Photoreceptor correlations

We know that cone sensitivities overlap, and that additionally, natural spectra stimulate cones types together, causing high correlations between the different cone types. L and M cone sensitivities overlap more than L and S or M and S sensitivities. Correlations are important for the model because our learning algorithm in V1 is based on Hebbian learning, which learns correlations in the input. Hebbian learning (in various alternative forms) is a widely assumed model of learning in the cortex (a model of synaptic plasticity; Shouval, 2007), so findings about the model are likely to be relevant for the cortex. If two channels are highly correlated, the AL model may not distinguish them (as for LISSOM; De Paula, 2007). However, we also know that the cone-opponent RGC channels should reduce these correlations (chapter 2). For example, responses of the RGC L-M and M-L pathways should be less correlated than the L and M cone responses are. Therefore, to discover what the impact of input channel correlations is on the model, and in the absence of suitable hyperspectral image data, our idealized model uses sRGB sensitivities for the IJK input channels instead of LMS sensitivities, since sRGB sensitivities are evenly distributed given the HSV-based hue manipulations described in section 5.3.1. We then add correlation back between the I and J channels in increasing amounts, to simulate increasing overlap. This dimension is controlled by d_C, with $d_C = 0.0$ meaning no additional correlation (idealized case), and $d_C = 1.0$ meaning I and J are identical (maximum additional correlation).

Figure 5.11 shows the effect of increased correlation between the I and J channels. While there is no impact on the overall cumulative photoreceptor and RGC activations, there is a clear impact on hue representation. V1 becomes increasingly unable to distinguish the I and J channels, leading to no distinction in preferences between hues that depend on varying amounts of I and J input. Additionally, the preferences segregate into only two types (K and I/J): two types of CR patch form, each representing only half the hues, and with no spatial separation between those hues. On the other hand, V1 continues to respond to all hues approximately equally overall.

The cone-opponent RGC is expected to reduce correlations (chapter 2), but because these RGC channels include a spatial component they remain highly correlated due to spatial (luminance) correlations. High correlation between two channels means the competitive Hebbian, correlation-learning AL model does not distinguish the channels. The same effect happens in the naive model because of the LMS correlations: the L and M cone sensitivities are highly correlated. This dimension contributes to problems **P1-range**, **P2-PCS**, and **P3-pref** because hues requiring I and J cone input separately cannot be represented distinctly, even though neurons will respond well to either one.

Figure 5.11 *(following page)*: *Dimension 3: photoreceptor correlation. Correlation d_C between the I and J channels is increased from left ($d_C = 0.0$; idealized model) to right ($d_C = 1.0$). Note: 0.0 does not represent no correlation, just no additional correlation beyond the correlation of red (I) and green (J) channels in the idealized dataset. Rows PR abs and RGC act show that cumulative photoreceptor and LGN activations are not affected, and the OR selectivity is also unaffected. However, increasing correlation between the I and J channels results in decreasing distinction between them, so that e.g. preferences that begin as separate red and green hues merge into yellow/orange. CR patches segregate into two distinct types. However, the cortex still responds approximately equally to all hues, albeit with responses to two ranges of hues that are spatially coincident. A question we will leave to the next chapter is why only two types of color patch result with high correlation between I and J, i.e. there no longer appears to be any combination of K and either I or J.*
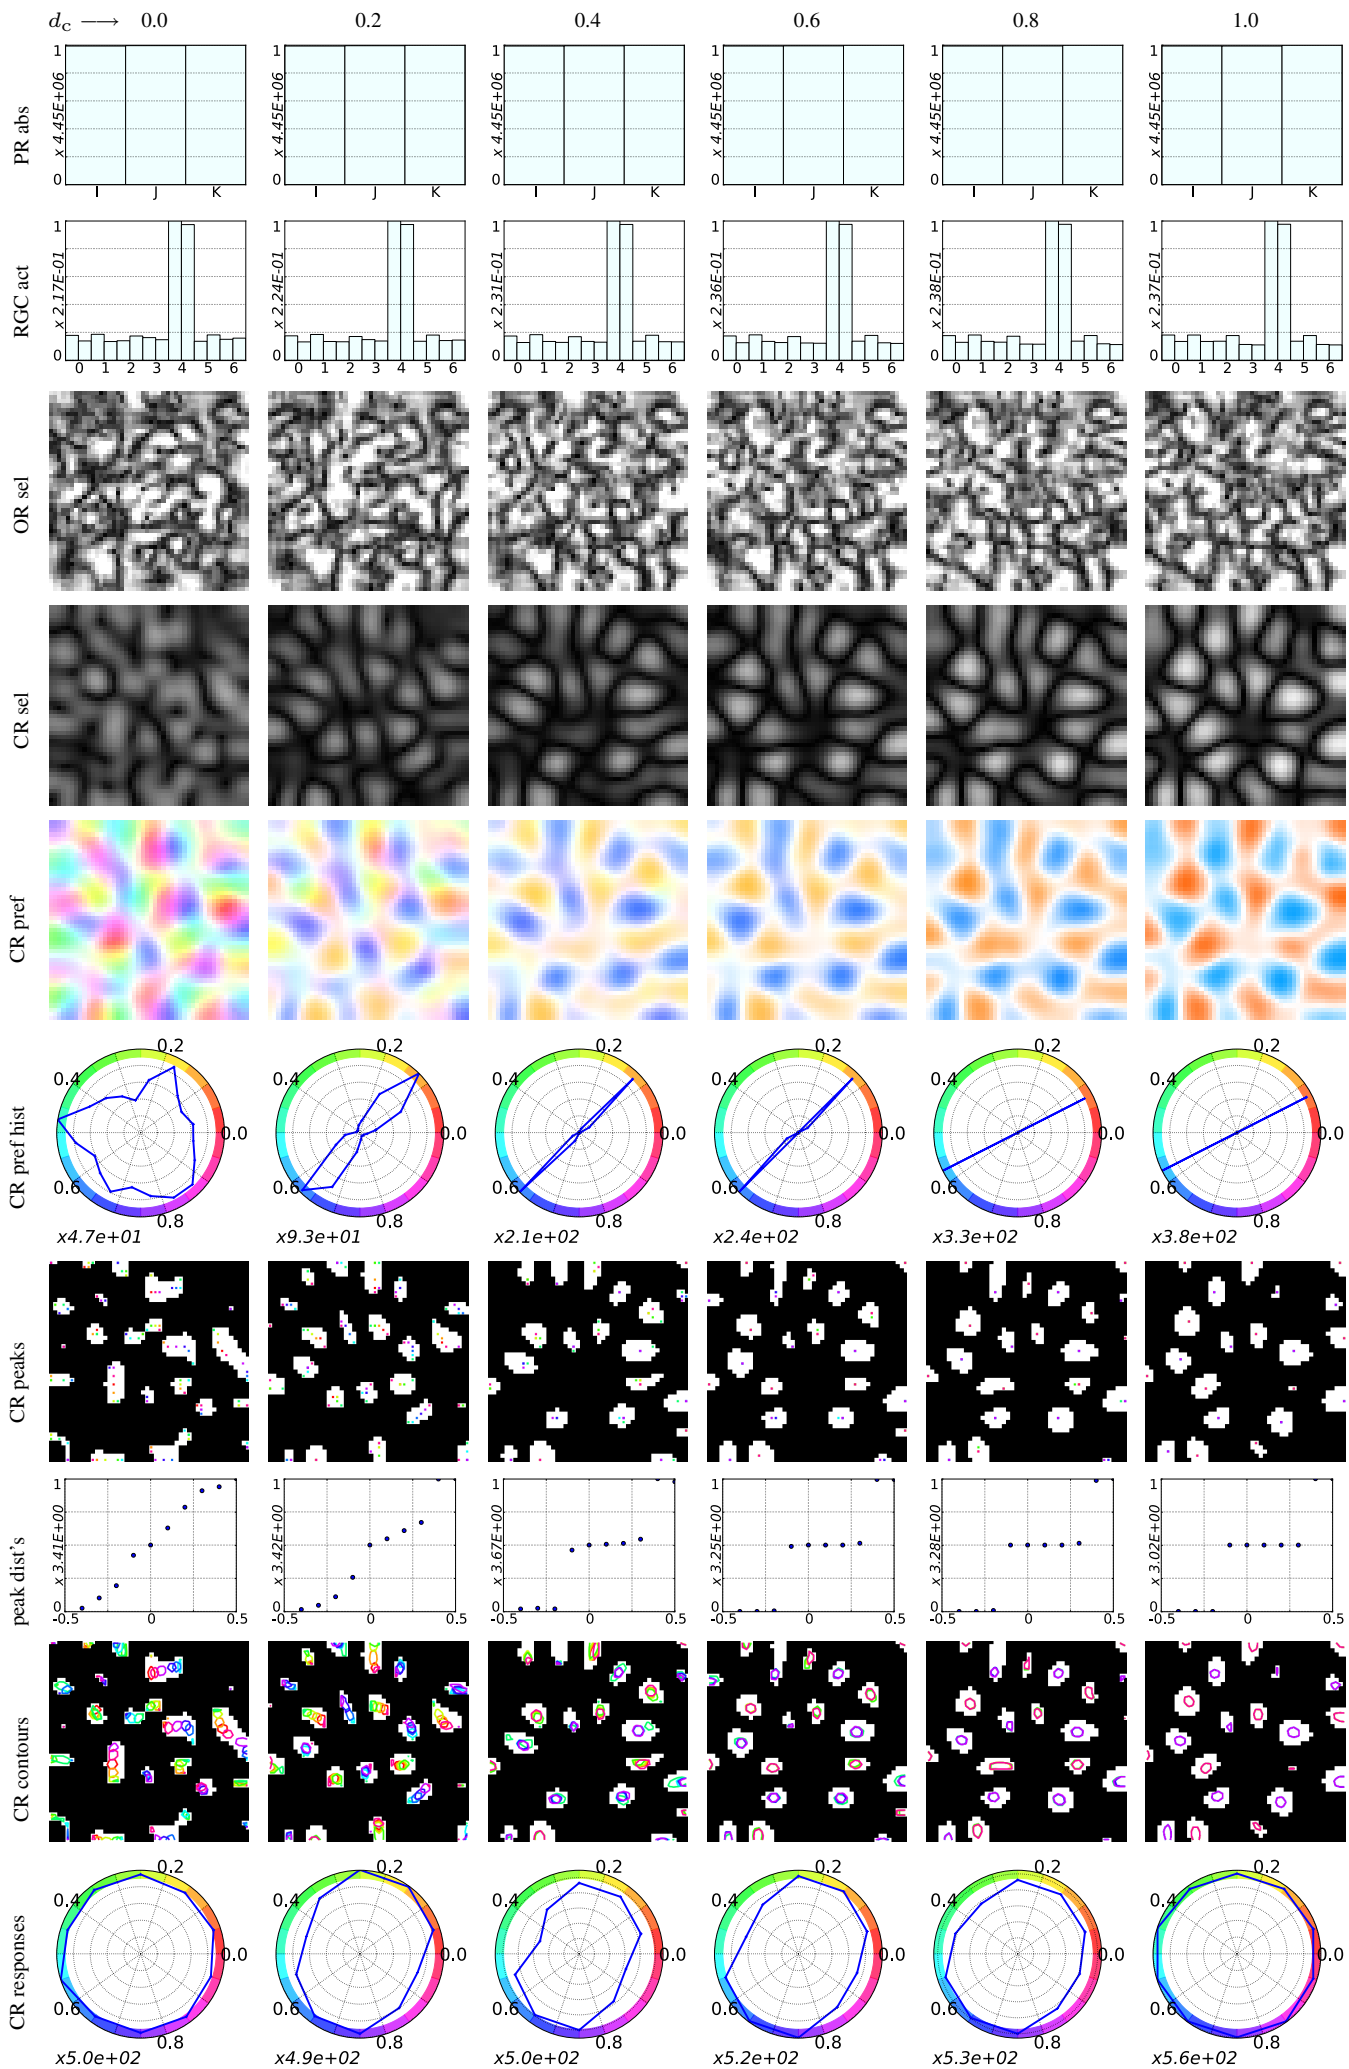
5.5.5 Dimension 4: RGC pathway symmetry

The naive model has an RGC pathway that does not combine (oppose) all cone classes symmetrically. The idealized model, on the other hand, has an RGC pathway that composes all cone classes symmetrically (figure 5.5). For this dimension, we consider the effects of these two pathways.

Figure 5.12 shows results from these two extreme positions on the RGC pathway symmetry dimension: $d_S = 1$ (idealized model, left) and $d_S = 0$ (right). While the symmetric idealized model develops a hue preference map containing all hues, the unsymmetric pathway results in a much more limited range of hues. Additionally, two distinct patch types occur: one a combination of I and K input ("red/blue"), the other a combination of J and K input ("green/blue"). Regions of the cortex preferring K alone (blue) are less selective, and not organized into patches. The preferences are reflected in the responses within CR patches: the symmetric pathway allows CR patches containing all or a wide range of hues to develop, with spatially separated peaks, but the asymmetric naive model develops CR patches that respond only to a limited range of hues (in two distinct groups, as indicated by the the peak distances row). Furthermore, the response plot row shows that the response to some hues (e.g. yellow) is lower than to others.

We can therefore conclude that any rescue of this dimension must at least balance the activities for the different hues, i.e. boost yellow. The results also indicate that having fewer discrete RGC channels makes it more difficult for intermediate hue preferences to arise. Again, this dimension evidently contributes to problems **P1-range**, **P2-PCS**, and **P3-pref**: the non-symmetric pathway causes uneven activities for different hues, prevents representation of full PCS developing, and leads to uneven hue preferences.

Figure 5.12 *(following page)*: *Dimension 4: RGC pathway symmetry. The non-symmetric RGC pathway (as used in the naive model) introduces a bias in color preferences and responses. In the right-hand column ($d_S = 1$), we can see that the preference map contains very few yellow-preferring units, and V1 has a lower overall response to yellow. There are only two blob types, and peaks are not well separated. On the other hand, in the left-hand column ($d_S = 0$; idealized model), the preference map is unbiased, and V1 responds equally to all hues. Using a coextensive "blue/yellow" (K vs. I+J) pathway does not qualitatively affect these conclusions.*
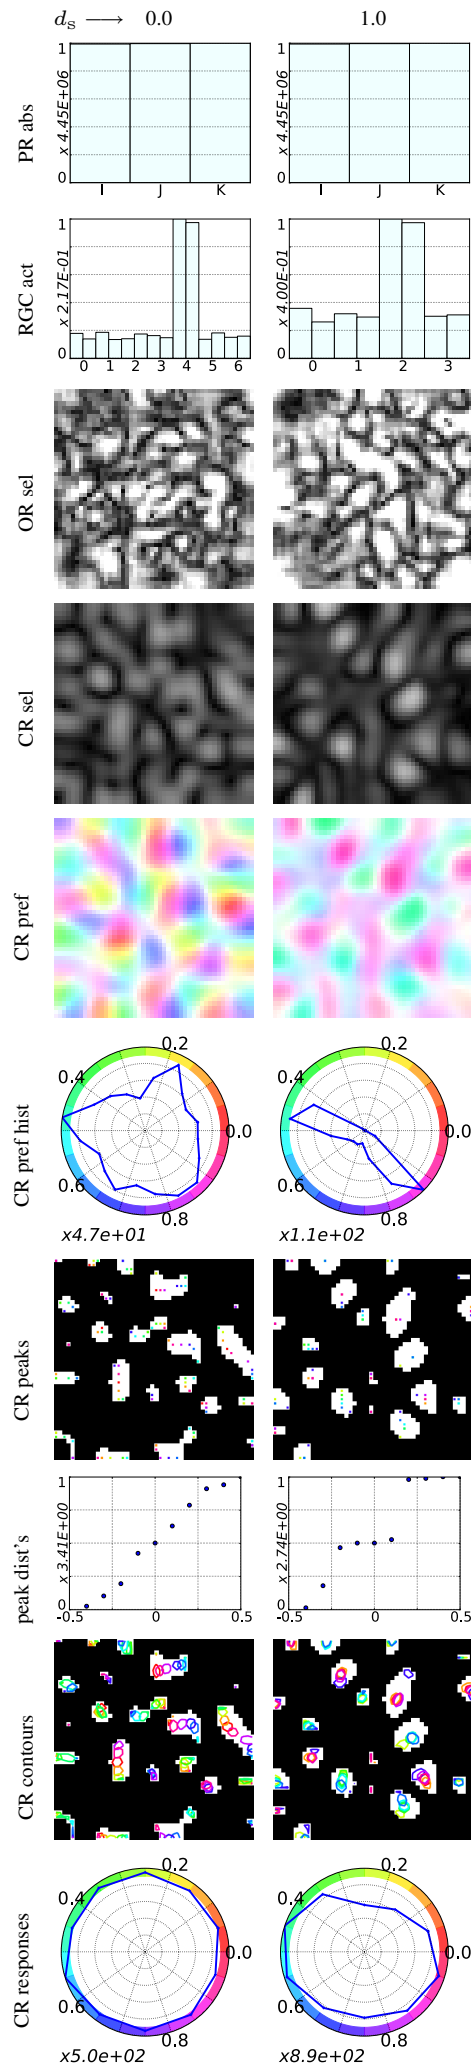
5.5.6 Dimension 5: LUM vs CR

We have seen that the idealized model's V1 forms a spatially contiguous OR map and has CR-selective units organized into isolated patches, matching results from macaque monkey. However, this particular balance of OR and CR maps is just one position along a dimension that ranges from completely OR dominated to completely CR dominated (figure 5.13). The position along this dimension is determined by parameter d_L, a measure of the fraction of total weights from V1 to RGC dedicated to LUM or to CR. For $d_L = 1.0$, the chromatic pathway has zero weight, for $d_L = 0.5$ the chromatic and luminance pathways have equal weight, and for $d_L = 0.0$ the luminance pathway has zero weight (though even in this case the cone-opponent channels carry luminance information too, since they are also spatially opponent).

Figure 5.13 shows that as d_L is decreased, OR selectivity decreases, while CR selectivity increases. OR selectivity never falls completely to zero, as expected. At low values of d_L, the CR representation is no longer patchy, and so although the CR map represents all hues more uniformly than for higher values of d_L, the representation does not match existing experimental results from macaques as well as at higher values of d_L.

This dimension helps to explain **ER2-patches** (CR in patches, rather than OR in patches), and is also relevant to **P1-range**: d_L needs to be low enough that CR patches can form, but if it is too low, CR will dominate and there will not be CR patches among a generally spatially contiguous OR map. Additionally, this dimension shows two interesting features of map organization, which we will discuss in the following chapter. Firstly, there is a balance of OR vs. CR selectivity, determined by the balance of the afferent CR and LUM signals to V1. However, we can also consider the spatial scales of the two maps: as soon as CR selectivity arises, it is the dominant feature in the map organization (i.e. it has a larger spatial scale). Previously, Miikkulainen et al. (2005) have shown, in a LISSOM model of the development of OR and DR selectivity, that the dominant feature depends on the speed of the moving input patterns. We will consider these observations futher in the next chapter, along with the correlation of CR and luminance in the input.

The position of this dimension used for the idealized model ($d_L = 0.6$) is chosen by matching the experimentally observed balance in macaque: CR is organized into patches, while OR maps are spatially contiguous. Different species could lie in different positions along this dimension, depending on their neural architecture. Different individuals could also lie in different positions along this dimension, e.g. because of their rearing environment or visual history. An individual raised in a monochrome

environment, for instance, may develop a V1 that ignores the inactive chromatic channels, and thus end up with an organization in V1 with no CR patches ($d_L = 0.0$), as would an individual who was born with only a single cone type.

Figure 5.13 *(following page)*: *Dimension 5: Balance of luminance vs. CR. V1 receives input from LUM and CR RGC channels. Adjusting the relative weighting of LUM and CR from $d_L = 1.0$ (left-hand column; no CR channels), though $d_L = 0.5$ (CR and LUM weights equal), to $d_L = 0.0$ (no LUM channels; right-hand column) shows the range of possible OR vs CR organizations in the model. OR dominates in the left-hand column, while CR dominates in the right-hand column. Note that even at $d_L = 0.0$ there is still an OR map because the chromatic pathways carry some luminance signal. Different species could appear in different positions on this dimension depending on their physiology, for instance if OR is de-emphasized, perhaps by having non-overlapping cone sensitivities, or no RGC LUM pathway (e.g. only very small RFs), they may appear towards the right-hand column. Different individuals could also appear at different positions, if they have sufficiently different visual histories (e.g. raised in highly colorful environment vs. raised in a monochrome environment). Existing experimental results from macaque are best matched by $d_L = 0.6$. It is interesting to note that a patchy organization of CR can still be found by thresholded selectivity or by differential response to RG-BW gratings even when the underlying CR selectivity is still reasonably high. This could help account for varying experimental findings (chapter 2) of OR and CR being represented separately or together. Another interesting thing to note, which we will discuss in the next chapter, is that the spatial scale of the CR map is always larger than that of the OR map.*
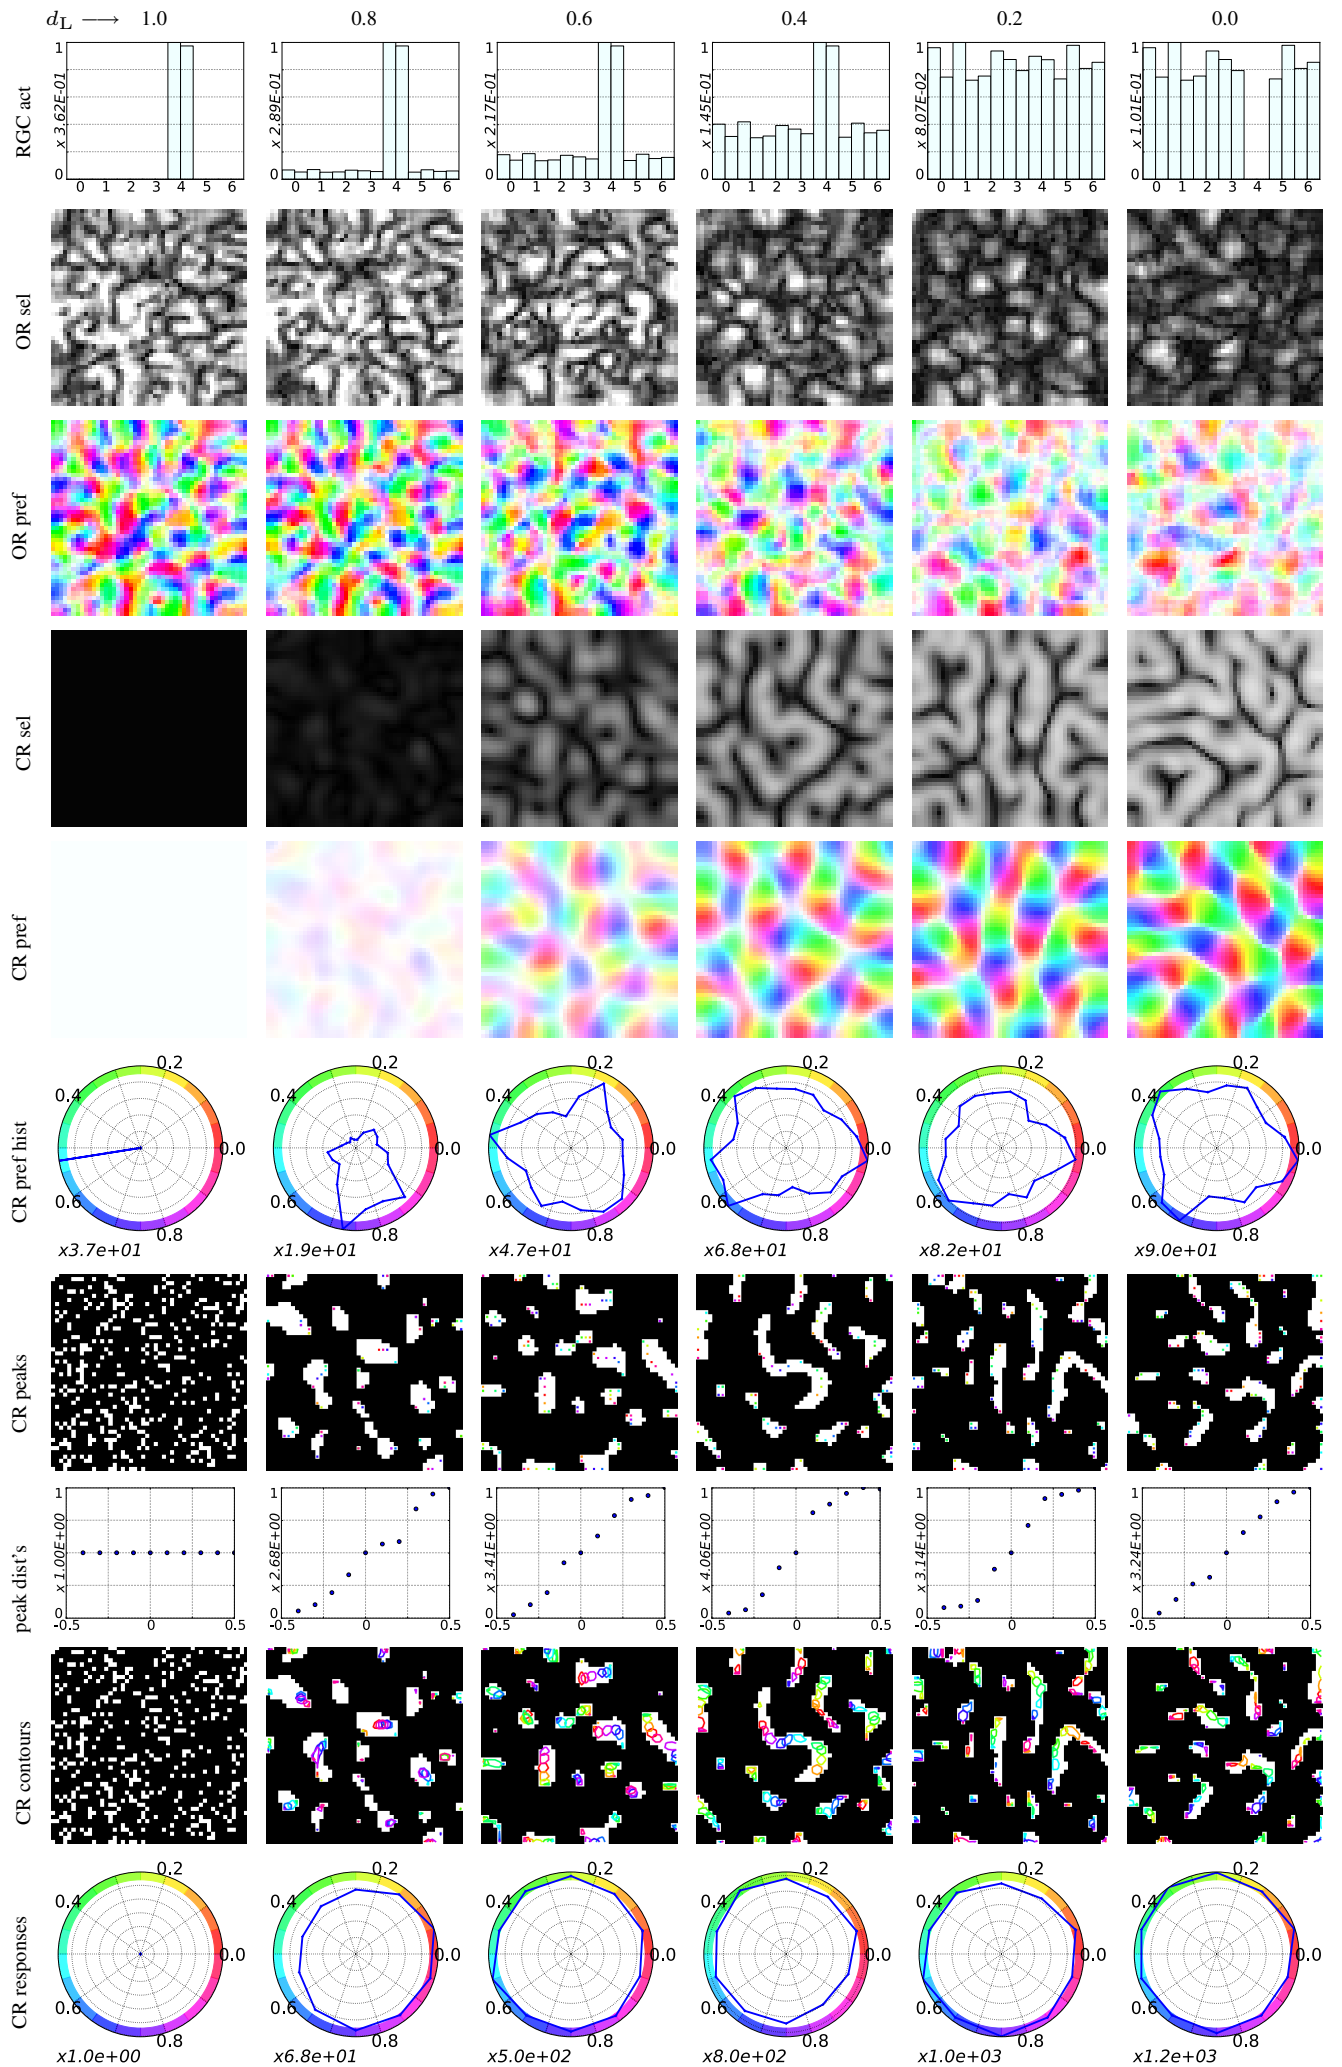
5.6   Discussion

The previous chapter began with experimental results showing OR-selective cells organized in a spatially contiguous map alongside CR-responsive cells in patches, and with cells in the CR patches organized according to PCS and hence able to code hue by spatial position. However, the model in the previous chapter—the naive model—could not explain the organization of hue according to PCS. In this chapter, we extended the naive model—leading to the idealized model—to show what is required for a realistic organization of hue to develop. For the first time, this gives a model with a hue organization matching experimental data, while simultaneously matching the overall organization of OR and CR.

In the process of creating the idealized model, we identified several important dimensions to which the the model is highly sensitive, but that the brain is apparently relatively insensitive to. Humans robustly develop a perceptual color space that appears to be reasonably consistent both between individuals and for one individual across time, despite very large changes in input caused by environment and physiology. The model, however, can only do this if we artificially eliminate these variabilities. These results are interesting because they indicate problems that the brain is overcoming in developing PCS, and so in the next chapter we will consider how the model might be made more insensitive to the dimensions. Below, we briefly review our findings about each of the dimensions, and afterward consider implications for experimental studies:

- Input hue bias: when the model is trained on images with a highly non-uniform sampling of PCS, which appears to be typical of visual experience, the model fails to develop a realistic organization of hue.

- Relative photoreceptor channel strengths: the model V1 requires the L, M, and S channels to be similarly active over the simulation's lifetime (or at least, long periods during development) in order that all channels are represented in the CR map. However, there are many factors that can prevent this happening: the environment could for example reflect more long-wavelength light than short-wavelength light to the eye on average; the eye could filter out more short-wavelength light than long-wavelength light before the light reaches the photoreceptors; the retina may have more long-wavelength than short-wavelength cones. These things may vary over time and space (field of view) for one individual, and between individuals, yet they appear to have little impact on human perception of color (chapter 2). The relative photoreceptor channel strengths dimension, which incorporates all these sources, illustrates that the model requires

the ability to adapt to uneven cone excitations. We demonstrated in the previous chapter that adaptive homeostasis in the photoreceptors could mitigate this situation.

- Photoreceptor correlation: the model V1 requires that L, M, and S channels are not too highly correlated, or else input channels cannot be distinguished and hues depending on distinction of those channels disappear from the CR representation. In reality, the L and M cones *are* highly correlated for typical visual experience given their overlapping sensitivity functions, so we believe there must be a mechanism to reduce this correlation. We expected that RGC cone opponency would do this, but it appears that the spatial element of the channels leaves too much correlation.

- RGC pathway symmetry: the naive model's non-symmetric RGC pathway appears to suppress some hues. A symmetric combination of cone-opponencies, as in the idealized model, allows all hues to be represented equally in V1.

- Luminance vs. CR: the model V1's balance of OR and CR maps is affected by the balance of LUM and chromatic channels, and a particular balance matches existing experimental results from macaque well. However, we expect that physiological and environmental factors could affect the balance of OR and CR in macaque V1, and may be different for other species depending on physiological and environmental factors.

As stated earlier, we believe the above are dimensions that the visual system is relatively insensitive to. Individuals grow up in different environments with different color statistics, have different cone ratios, have lenses that yellow over time, etc., and yet the brain is fairly insensitive to those differences—PCS seems to be fairly consistent between people and over time. It is surprising that the color vision system is not more sensitive to some of these, because we know that both the environment and physiology have a major effect on color vision development (chapter 2). The variabilities *are* reflected in the visual system, for instance the different statistics of different color environments do affect PCS, or different cone sensitivities do cause different detailed color matches between individuals (chapter 2), but the effects are much smaller than we might naively expect.

5.6.1   Accounting for the range of experimental results

Beyond developing a representation of hue that matches experimental data, and learning what problems the model needs to overcome to do so, another interesting consideration is what the model can tell us about interpreting the existing experimental results.

As reviewed in chapter 2, there is disagreement in the literature about several aspects of the representation of hue in V1. We have focused on data from optical imaging (OI), backed up by preliminary in vivo two-photon calcium microscopy (2P) studies, because imaging allows cells to be simultaneously imaged over a wide enough area, yet at a high enough resolution. However, some studies using other methods such as electrophysiology have reported e.g. one "color opponency" per patch (i.e. responses to RG-BW gratings separate from responses to BY-BW gratings), rather than multiple hues per patch. In our model, we saw that even while varying only the input seed for hue jittering, when we determined peak responses in the most selective patches, we sometimes found patches representing only a limited range of PCS. Despite this, the underlying hue representation was overall the same in all cases.

Similarly to the question of hues represented in each patch, experimental studies have in some cases found that OR and CR are represented separately, while in other cases are represented together (chapter 2). Our results indicate that even when the underlying CR selectivity is quite uniform, a technique that defines patches by the differential response to RG-BW gratings, or by thresholded selectivity, can find a patchy distribution of CR. This indicates that experimental conclusions may differ for a common underlying reality.

Furthermore, experimentally imaged subjects could have been raised differently, so could be in slightly different positions along dimensions. For instance, individuals may vary along the luminance vs. CR dimension, and this could account for some studies finding that color-preferring cells and luminance-preferring cells are quite overlapping (i.e. that cells represent both features), while others find the two types of cell are separate. We do not know if the different experimental findings are caused by genuine differences between individuals, or from experimental techniques. Experimental results are limited to only a few individual macaques so far. We will discuss in chapter 6 how the model could be used to help predict results of experiments not yet done, such as imaging individuals reared in environments with different color statistics.

So far, hue organization has only been investigated experimentally using OI and 2P for a few individuals (with unknown visual history), leading to a limited amount of data to

compare against. Performing more quantitative analysis beyond comparing distances between peak responses to different hues is therefore currently of limited value. However, once results from more individuals (with more information about their visual history) are available, it will be possible to perform more quantitative comparisons between experimental and model results.

Finally, one more issue with experimental work that is apparent from this chapter's modeling work is the difficulty of controlling "color" stimuli appropriately. Stimuli designed to test hue preference, but that possibly vary in other aspects of PCS (e.g. having a different saturation), make it difficult to determine hue preference with certainty. No experimental studies have "integrated out" other features of PCS (e.g. presenting patterns with multiple saturations for each hue, along with multiple values of other features such as orientation and direction of motion), presumably partly because of the impractical amount of time that would require, and partly because the true dimensions of perceptual color space are unknown. However, an important aspect of successfully coding for hue is to have a stable representation, as has been found for orientation, so we will return to this issue in chapter 6.

5.7   Conclusion

Experimental work indicates that CR selective cells are organized into spatially segregated patches, and that within each patch a wide range of hues is represented, organized according to PCS. In this chapter, we have demonstrated that the idealized model can develop a realistic organization of hues within CR patches, alongside a contiguous OR map. This is the first time all these aspects of the experimental data have been matched. To achieve this, the model must artificially be made relatively insensitive to a number of sources of variation, including input color statistics and layout and numbers of photoreceptor types. We have shown the effect of each of these, and thus we can see what modifications (rescues) would need to achieve for a model to develop a realistic representation of PCS. This will guide future work to implement rescues via biologically plausible mechanisms.

Chapter 6

Discussion and future work

In this chapter, we will review the work presented in this thesis, and consider how successful we have been in meeting the goals set out in chapter 1. Afterward, we will catalog the suggestions for future experimental and modeling work that have arisen throughout the work in this thesis.

6.1   Evaluation of goals

Here we will review findings from the preceding chapters, and consider whether we have reached the goals set out in section 1.4:

- In chapters 1 and 2, we reviewed previous work showing how the perceived color and physical stimulus are not directly related. Although all aspects of vision can illustrate how perception and reality differ (e.g. the perceived orientation of a line can differ from its physical orientation, or motion can be perceived where there is none), the difference between physical stimulus and perception is particularly clear for color. We also saw evidence that the representation of color by the early visual system during development is both highly affected by the environment and underlying physiology, yet at the same time the development of perceptual color space (PCS) (as judged by experiments testing visual performance) appears to be affected surprisingly little by significant variabilities in both. Excitingly, we then saw that there appears to be an organization for perceptual color in macaque primary visual cortex (V1), which can be measured and related to concrete neural anatomy and physiology. *Goal 1* was to show that the neural representation of color is a useful model for understanding how the brain

represents and processes information in general. By outlining how the physical stimulus and perception differ, but at the same time showing there is a real possibility of imaging the neural representation of color in the brain, we set up *goal 1*. But how could the necessary wiring arise? Experimental studies so far do not address this. Turning to computational models, we found that existing models of the development of orientation (OR) and color (CR) in cortex are not able to explain important results about the organization of PCS, in particular how hue is represented. Therefore, chapter 2's investigation demonstrated the need both for new modeling work, and to better understand visual input in terms of statistics relevant to color. These were the subjects of the following chapters.

- Chapter 3 showed that (a) the distribution of colors in natural scenes is a highly non-uniform sampling of PCS, with natural scenes being dominated by certain colors such as yellow and green, and (b) the distribution of cone absorptions is also biased in several ways that are problematic for the type of models we reviewed in the previous chapter. We also reviewed databases of calibrated color images suitable for use as input to the model, and characterized two specific databases. This left us with concrete input to use for an input-driven developmental model, and helped to understand how that input would affect the model. For us, this investigation was a necessary step on the way to next chapter, but is worthy of future study in its own right.

- *Goal 2* was to integrate biological findings about natural images and the physiology of the early visual system into a simple computational model, to see whether such a model could develop a realistic perceptual color space. Hence, chapter 4 began by solving a number of problems not addressed by state of the art models of the development of orientation and color processing, showing that adding adaptive mechanisms at the photoreceptor and V1 levels allows a model to develop orientation and color maps reliably. To determine whether the color and orientation maps developed by the new model (the naive model) were realistic, we made a detailed comparison with experimental work. We found that the naive model has better results than any of the previous models, while being true to known anatomy. However, as for previous models, we found the naive model did not develop a realistic representation of hue, crucially being unable to support a spatial coding of hue as is apparently found in macaque monkey V1. Therefore, to achieve *goal 2*, we needed to analyze the naive model in light of known biology from chapter 2.

- As suggested above, by analyzing the model of chapter 4, in chapter 5 we were

able to create a new model that does develop a realistic hue organization, matching in detail recent experimental results from macaque and achieving *goal 2*. We showed that previous models are unable to adapt to variabilities or bias in visual input and physiology, and hence are unable to develop a realistic neural representation of perceptual color space. From these findings, we hypothesized that adaptive mechanisms are critical to the development of color vision, allowing the visual system to be relatively insensitive to sources of bias and variability. Generating such a hypothesis was *goal 3*. Mechanisms such as homeostasis and adaptation are well accepted biologically, but typically omitted from models of luminance and color processing. As stated above, we think the neural representation of color is a useful model for understanding how the brain represents and processes information in general, precisely because such adaptive mechanisms cannot be omitted from even the most basic models of color processing if they are to be successful. We therefore hope we have achieved *goal 1*.

The final goal of this thesis was that our work should suggest future experimental investigations, and lead into new modeling work, which together can be used to evaluate our hypotheses. During this thesis, we have encountered a number of problems evaluating and interpreting experimental results, so we can therefore suggest experiments that would help resolve these issues. Secondly, the concrete result of this thesis is two novel models that form a platform suitable for further investigation of the adult color visual system and its development. There are a number of interesting experiments that could be performed on the models to compare with known biological phenomena. The results of such experiments could suggest mechanisms for the biological phenomena, or introduce new constraints on the model, guiding improvements to it. Finally, the model has a number of concrete shortcomings and omissions; solving any of these would extend the valid range of the model or open up new biological phenomena for investigation.

6.2   Proposed experimental work

In order to resolve some of the questions about how PCS is represented in macaque primary cortex, below we describe three specific themes that have emerged from our modeling work. To briefly motivate the themes, we can give an example. Recent interesting functional magnetic resonance imaging (fMRI) studies have found the representation of hue in V1 is not uniform. Hue representation was found to biased along a lime–magenta axis in human V1 (Goddard et al., 2010), but separately, hue represen-

tation was found to be biased along an orange–cyan axis in macaque V1 (Lafer-Sousa et al., 2012). The difference could be related to all or a combination of: different anatomy between species; different prior visual experiences; differences in the color properties of the stimuli used during testing; or difficulties from averaging V1 color response at the resolution of fMRI. For our modeling work, chronic imaging in vivo, at two-photon calcium microscopy (2P) resolution of individuals from the same species, reared in controlled environments of various known color statistics, with test stimuli whose physical and perceptual color properties are accurately known, would be the best match, and therefore allow us to resolve ambiguities in the model. However, each of these requirements presents significant technical challenges on its own, and achieving all of them is unlikely to happen in the near future. Therefore the rest of this section focuses on specific, more tractable goals.

6.2.1   Hue preference maps

Currently, there are no published hue preference maps for macaque V1 (as measured by optical imaging or 2P). For orientation, it has been possible for example to present combinations of sine gratings of different orientations, and record neurons' preferred orientations (i.e. the orientation that causes the strongest response). For hue, only maps of the peak location in response to each hue have been published. It would be useful and interesting to have the underlying responses, in order to compute a hue preference map. Our modeling suggests apparently uniform hue peak maps can be obtained from quite non-uniform preference maps. Experimental hue preference maps measured using optical or in vivo two-photon calcium microscopy (2P) imaging, even if appearing to show no interesting features (such as showing only one preference, or showing preferences that are not stable) would help constrain the model, which in turn would allow us to improve our understanding of how the maps develop.

A potentially interesting application of the model in this respect could be to test proposed experimental procedures before they are performed. This could provide valuable insight into what information a given experiment could reveal under ideal, controlled conditions. For instance, in the model, measuring a hue preference map with stimuli that vary in hue but also incidentally in saturation can result in a misleading preference map—but it is easy to realize that mistake because, unlike in an animal, everything can be inspected. A close link between experiment and modeling in this way would also greatly benefit the modeling work, because essential features of the experimental procedure could then be replicated. The model could also serve as a method of recording

and communicating the procedure.

Versions of our model with no adaptation illustrate clearly that representing perceptual hue is different from representing wavelength. Although features such as the spatial organization of peak responses to different hues defined with stimuli that also vary in wavelength are important (e.g. is "red" close to "magenta" as well as being close to "orange"?), the next step is to ask whether such a map genuinely represents perceptual color, or is closer to physical wavelength (Zeki, 1983). With no short-term, spatial adaptation in the model, simulating a change to the ambient lighting conditions under which a hue map is measured would result in a different hue map, illustrating the difference between perceptual hue and wavelength preference. Hue maps in animals could therefore be measured with stimuli that vary perceptually but not physically, or vice versa, e.g. by presenting more complex patterns as has been done only for single cells (Zeki, 1983), or by using temporal adaptation.

Finally, while there have been several studies of OR preference using 2P, confirming the detailed organization of OR preference maps, only preliminary results have been shown so far for CR. The single-cell resolution of 2P, combined with being able to measure a large enough area, makes it an ideal match for our modeling work. If the organization of hue in V1 is largely contained in small patches, organized internally on a finer scale, 2P imaging will be necessary for more certainty about this organization. Of course, the same difficulties of selecting the appropriate stimuli to use will exist as for optical imaging or any other method.

6.2.2 Development of hue representation over time

We know how orientation maps develop over time from birth, and that maps exist pre-natally (established by e.g. retinal waves). But we do not know anything about how CR organization develops, or whether any form of CR map is present at all before birth. L and M cells are genetically very similar, and even adult brains can adapt to new cone types (Mancuso et al., 2009), leading us to assume that a prenatal, entirely experience-independent explanation of hue preference maps is unlikely (though inter-action between the two modes seems likely). Any data about the time course of hue representation's development would therefore be extremely valuable for constraining the model.

However, our modeling work indicates that it is critical such data be combined with knowledge of the color statistics of the rearing environment. Ideally, chronic map measurement would be combined with control over the visual experience during de-

velopment, though this may be difficult. But it is at least already possible to perform chronic imaging studies over long periods in behaving animals (Arieli et al., 2002). For example, Li et al. (2014) imaged macaque monkey V4 a month apart to confirm stability of hue peak maps. Combining such long-term imaging with controlled color environments would also be interesting, because we know adult color vision is highly adaptable over long periods (weeks, e.g. Neitz et al. (2002), and years, e.g. Werner et al. (2004); see chapter 2). Alternatively, such imaging could be combined with e.g. adding a new photoreceptor pigment to dichromatic animals (Mancuso et al., 2009) and imaging V1 over time, while also monitoring progress in retinal expression of the photopigment.

6.2.3   Determine what is realistic visual experience

Ideally, we would like to collect calibrated color images from the rearing environments of the monkeys imaged in existing studies. However, the environments are unknown (chapter 3). If the environment was not the natural one we assumed, but was instead one with many colorful man-made objects, this could have implications for the hue maps being measured, as discussed above. Determining realistic visual experience is interesting not only for interpreting experimental results about hue representation, but also in its own right. Chapter 3 indicates we are exposed to a highly non-uniform sampling of our perceptual color space in nature, with little experience of certain colors. Is this genuinely the case? Chapter 3 indicates that taking more images in similar locations would not change the conclusion about the types of scenes sampled in that chapter, but collecting calibrated color image databases from different regions (e.g. tropical jungle) or at different times of year could reveal interesting differences. Additionally, rather than manually photographing particular objects or scenes, it could be possible to attach calibrated RGB cameras to animals. Results from collecting such image databases could be analyzed as in chapter 3, or possibly using more advanced color appearance models.

Beyond characterizing visual experience in terms of light entering the eye, it may also be necessary to consider the significance of different colors to an organism. For instance, the colors of a venomous snake might be experienced less frequently than the colors of grass, but the snake's colors are likely to be more significant and hence could have a greater impact on learning mechanisms. Therefore, it may additionally be necessary to take salience into account when determining how the environment affects hue representation in the cortex.

6.3 Proposed modeling work

Constraining models like the ones we have presented in this thesis is not easy. One source of constraints is results about adult color vision performance (i.e. from psychophysical studies), which constrains how the model should perform after development. These constraints could reveal mechanisms that need to be changed or added to the model. Or, the model can be used to explain experimental results, or to predict new ones. However, a large amount of such data already exists in the psychophysics literature; unifying the data into a more general model that is subject to many other constraints forces us to evaluate the consistency of models of individual, separate phenomena.

We have contributed a computational model that can act as a platform on which many experiments could be performed, and we list a selection of immediately accessible ones below. The models used in this thesis are included with the Topographica simulator to allow the future work proposed below to be carried out, and also to enable reproducibility (an important but surprisingly difficult aspect of computational neuroscience; Crook et al., 2013), and to allow the model to be incorporated with other aspects of cortical processing (Bednar, 2012). One person or small team can typically only include in a model the features necessary to support the specific phenomenon they are studying, but this can result in a large number of specific and potentially incompatible models. Making models easily accessible, and attempting to integrate them, is a worthwhile endeavor.

While the models as included with Topographica are the best place to start for new work, for completeness all results figures in this thesis may be regenerated automatically by Makefile (including running the underlying simulations) from the freely available source (cb2014thesis.sourceforge.net). With the advent of tools such as IPython notebook, this kind of approach is becoming both simpler and more powerful (Stevens et al., 2013a).

6.3.1 McCollough Effect

Figure 6.1 shows the McCollough Effect, an aftereffect in which the perceived color depends on orientation. Previous work has modeled the effect in a dichromatic (red/green) model using LISSOM, and has shown that the change in perceived color could be explained by adaptation of lateral connections in V1 (Ciroux, 2005). The model presented in this thesis would allow the effect to be modeled for different pairs of colors

(Stromeyer, 1969). Preliminary work to do this has begun (Spigler, 2014), but indicates that the balance of orientation and color in the model is an important factor in being able to reproduce the effect, therefore potentially helping to constrain the representation of OR vs CR (see section 6.4.4 below). Future work could also consider the effect's dependence on wavelength instead of perceived color (Thompson and Latchford, 1986), to help place constraints on photoreceptor adaptation (e.g. adding spatial adaptation; section 6.4.7), or on the effect's transfer between eyes (e.g. if ocular dominance were included in the model; see section 6.4.8). There are also related orientation/color effects, such as the tilt aftereffect's dependence on color, and others (see e.g. Clifford et al., 2003).

6.3.2 Angular dependence of color aftereffects

We saw in chapter 1 that perceived color depends on previously viewed colors. Webster and Mollon (1991) found that perceived color changes in a way that depends on the difference in color between adapting and test stimuli (figure 6.2). Such an effect could be tested in the model. If there is a mismatch, the human data could then be used to constrain the model. Or, if the effect is found in the model, the mechanism of the effect in the model could be used to hypothesize about the cause of the effect in humans—similarly to the McCollough affect above, or to previous investigation of the tilt aftereffect in a related model of orientation (Bednar and Miikkulainen, 2000).

(a) (b)
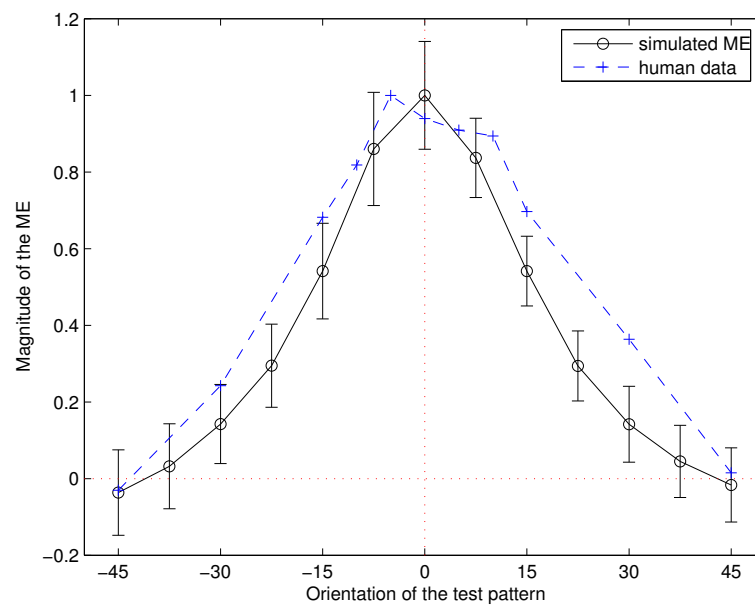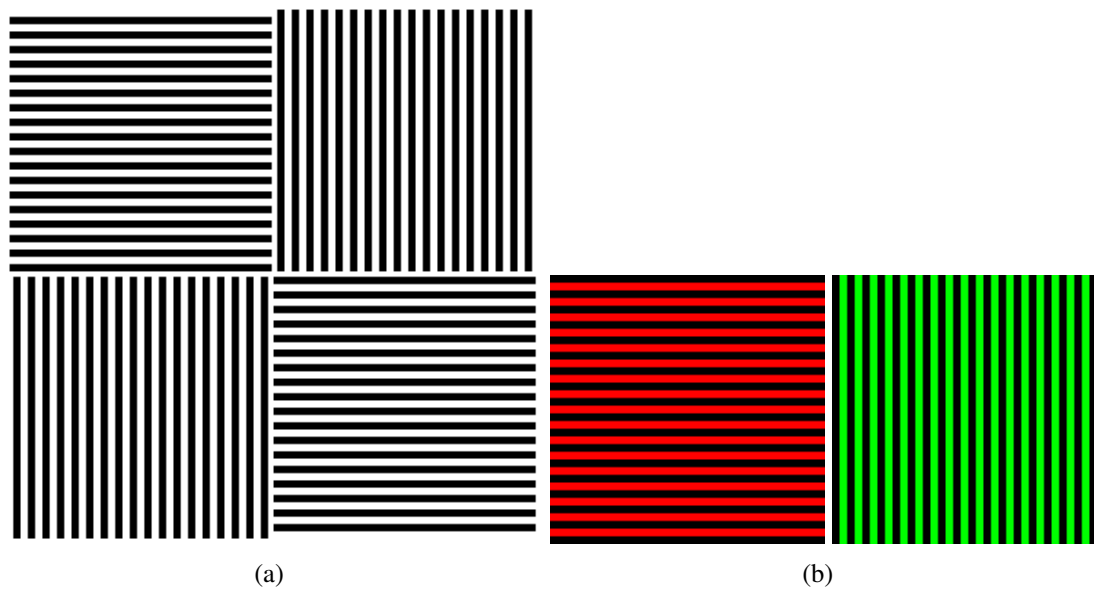
(c)

Figure 6.1: *The McCollough Effect. (a) Initially observe no colors on the achromatic (test) gratings. (b) After alternately gazing at the two chromatic (adapting) gratings for a few minutes (avoiding afterimages by moving your eyes around slightly), you should then observe that the horizontal grating is greenish and the vertical is reddish. The more time spent gazing at the chromatic gratings, the stronger the effect will be. Try adapting for five minutes and check back in an hour to see that the effect lasts. (c) The effect is maximal when the adapting and test gratings are orthogonal, while at $45°$ there is no effect (when you have the effect working, rotate the paper to see the color go away, and keep rotating to see the opposite color return!). In a dichromatic model, previous work has found that the angular dependency can be reproduced in the model, largely by adaptation of lateral connections in V1 (Ciroux, 2005). (c) reprinted from Ciroux (2005); data from Ellis (1977).*
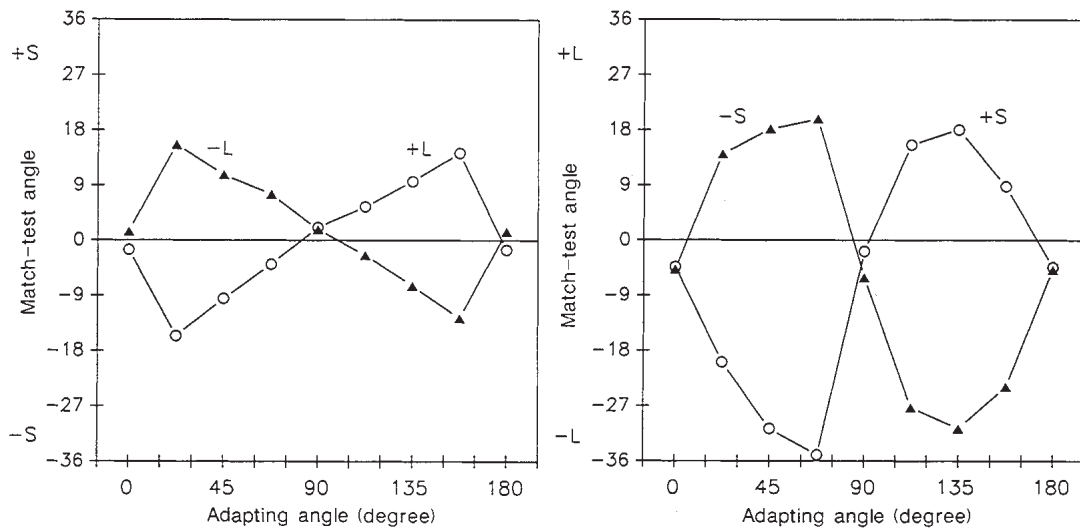
Figure 6.2: *Angular dependence of color aftereffects. Viewing a colored adapting pattern changes the perceived color of a subsequently viewed pattern. The change in perceived color varies with the difference in color between the adapting and test patterns (here the color space is DKL). Such data could be used to constrain a model's color processing. Reprinted from Webster and Mollon (1991).*

6.3.3 Representation of other aspects of PCS

So far, we have considered only the representation of hue. However, even in the simple perceptual color space used to analyze models in this thesis (HSV), there are two other dimensions—saturation and value. The hue representations we found are stable across varying saturations, for instance, but we did not explicitly test for an organization for saturation. Although no comparable data is available for macaque V1, hue representation in V4 appears to be stable across saturations, and there may be organization for saturation (in HSL space; Li et al., 2014). Additionally, in section 6.4.7 below, we discuss adding spatial adaptation to the model. Progress in this direction would allow investigation using colorimetrically well-defined spaces that take into account the background, such as CIE L* C* h_{ab} color space (CIELCh$_{ab}$) (which has lightness and chroma, and a more perceptually uniform representation of hue).

6.3.4 Tuning of neurons

What is the color tuning width of neurons in the model? By tuning width, we mean a comparison of a neuron's responses to stimuli of varying hue angle. A highly selective neuron would have a narrow tuning width, responding significantly more to a small range of hue angles than to the rest. An unselective neuron would respond the same for all hues. Two aspects of tuning width could be important for other future work suggested in this chapter. First, the tuning width could determine whether or not the angular dependence of aftereffects such as the ME (section 6.3.1) matches human data. Second, investigating how stable the tuning width is across changes in other dimensions of perceptual color space will show how reliable the hue representation is for coding hue (e.g. like contrast-invariant–tuning is important for orientation coding; Antolik, 2010; Stevens et al., 2013b).

6.3.5 Training from calibrated databases of different environments

As more collections of calibrated natural image databases become available, it will be possible to train the model on these and, if the environments are sufficiently different, to demonstrate the impact of different environments on details of perceptual color space. However, there are already images of urban scenes in both the McGill Calibrated Image Database (MCID; Olmos and Kingdom, 2004) and Barcelona Calibrated Image Database (BCID; Párraga et al., 2010), and there are more hyperspectral images of urban indoor/outdoor scenes than of natural scenes (e.g. Chakrabarti and Zickler,

2011). Using such images to train the model could already generate interesting hypotheses about the impact of environmental color statistics on the representation of PCS in V1 for modern humans, or animals raised in man-made environments vs. being wild caught. Additionally, as discussed in chapter 3, where an image database allows the recovery of reflectances, different illuminants could be simulated (for time of day, or weather conditions, for instance).

6.4 Improving the models

In chapter 5, we showed a number of dimensions representing physiological and environmental variabilities and biases to which the naive model must be made insensitive in order to represent hue in a realistic way. We demonstrated *what* needs to happen in the model in order for a realistic hue representation to develop alongside a realistic organization of orientation and color selectivity, which, along with the biological background reviewed in chapter 2, allows us to hypothesize about *how* it could happen in the model and hence the brain. However, our hypotheses will need to be tested via concrete implementations, which will require quite extensive work in each case. However, we have laid out what the problems are, and have provided a platform from which to test and learn about their effects individually, which will be useful even if our hypotheses about how to address the problems turn out to be wrong.

Our first hypothesis is that adaptation and homeostasis at each level of the model is critical for the development of perceptual color space. Our second hypothesis is that an additional mechanism to reduce input channel correlation is required. In the following two sections, we discuss these hypotheses and how they relate to the models of chapter 4 and 5.

6.4.1 Adaptation and homeostasis

The relative photoreceptor channel strengths dimension showed the effect of an input channel being activated more or less than others over long periods. Such an input channel imbalance could arise because of pre-receptoral filtering by the lens, different ratios of cone types, reflected spectra containing more energy over some wavelength ranges than others, or, in the extreme, missing photopigments or non-functional receptor types. The naive model showed that some form of photoreceptor homeostasis is necessary. The next step is to implement homeostasis via a realistic mechanism, as has

been done for the cortical component of the model (Stevens et al., 2013b).

Second, we think the same kind of mechanism will be required at the RGC stage. Firstly, because of photoreceptor correlation, even though the activities may be equal, differencing correlated vs. uncorrelated activities leads to imbalances in activity in the RGC channels. Secondly, we saw in the RGC pathway symmetry dimension that the unsymmetric pathway of the naive model leads to suppression of some hues; some balancing of the separate pathways may be necessary. Thirdly, as we will discuss in section 6.4.3 below, a more realistic subcortical architecture ("random" wiring of the RGC) is likely to require some form of RGC adaptation/homeostasis.

Note that a number of biological results could well be explained (or partly explained) by such mechanisms (and where the effects are not explained, will provide data to constrain the mechanisms). As reviewed in chapter 2, these results include deliberate experiments such as "red" goggle wearing (e.g. Neitz et al., 2002), as well as observed phenomena such as the lens yellowing with age, or a photopigment not being present from birth.

6.4.2 Correlation

The correlation dimension in chapter 5 demonstrates that as two of the input channels are increasingly correlated, V1 fails to distinguish them. In the model, this correlation can arise because of the overlapping cone sensitivities, and also from additional correlations caused by natural images. We expected that the RGC and LGN processing would reduce photoreceptor correlations, as the processing opposes input from different cone classes (Zaidi, 1997), but we found that correlation remains high because of strong spatial correlations in the images. Therefore, a modification of the model, or introduction of new processing, is likely to be necessary to solve this problem; we suggest some possibilities below.

Firstly, in biology, photoreceptor output to the retinal ganglion cells may in effect be less correlated than the initial correlations (i.e. the overlapping LMS sensitivities and their additional correlated stimulation by natural images). For instance, this could be because of lateral inhibition between cones via horizontal cells, or nonlinearities in photoreceptor responses (Baylor et al., 1987; Endeman and Kamermans, 2010) that may reduce correlations.

Alternatively, correlation may be reduced at the RGC stage in the model by the same kinds of mechanisms as above. Inhibition between ganglion cell outputs at some

point before V1 is not supported by current evidence (Solomon and Lennie, 2005), but nonlinearities in RGC processing have been shown to significantly reduce correlation (Pitkow and Meister, 2012). Correlation could also be reduced at the RGC stage if the chromatic RFs were smaller, reducing spatial correlations. An additional possibility is that simulating a more realistic cone mosaic, and more realistic (less perfectly overlapping) RGC surrounds, could reduce correlation (see section 6.4.3 below).

Potentially, the solution to this correlation problem could be constrained by results about color vision performance of individuals who have even more similar L and M pigment varieties than is typically found in humans (chapter 2). For some degree of similarity, the cone responses evidently cannot be distinguished.

Another aspect of correlation to consider is related to how similar hues are represented close together in the model cortex. In the model with no adaptation, colors (i.e. receptor activation ratios) that cause similar activation patterns will typically be represented close together (because of short-range excitatory connections). L and M cone sensitivities have much higher correlation than L and S cone sensitivities, so with no other mechanism, the L and M cones tend to be represented together, and separately from S cones. Natural images increase correlations, particularly in L and S cones, which we might expect to cause S and L cones to be represented together too, but in fact we found in the naive model that S cones are generally still represented in separate patches from L and M cones. Correlations must therefore be low enough that the cone types can be distinguished, but also the correlations between pairs of receptor types may need to be similar for them all to be represented nearby in V1.

Along with investigating what causes CR patches to form, and what causes the "hue circle", it would be interesting to study what influences whether hue peak maps within patches defined by RG-BW gratings, or based on selectivity, are found to join up or not. That is, whether peaks for red and blue are as close as those for e.g. red and green, and separated by colors such as magenta. As in experimental data, peak separations in the model are measured within restricted patches (e.g. those defined by RG-BW gratings, or by taking the most selective regions); this can result in patches disguising the underlying proximity of hues, and could account for some studies finding different patches representing different hues, and other studies finding many or all hues within a patch.

6.4.3   More realistic subcortical model

As mentioned above, the ratio of L and M photoreceptors varies greatly between individuals, and there are relatively few S cones (chapter 2). Our model in chapter 5, the idealized model, shows the result of imbalanced channel strengths, and it should be straightforward to simulate e.g. a realistic S cone density (approximately 5 %) or the frequently mentioned 2:1 ratio of L:M cones in human (though not enough data has been collected to know how accurate these estimates are; chapter 2). The effectiveness of photoreceptor homeostasis could then be evaluated on concrete biological variabilities. One other interesting possibility might be to compare the organization of OR and CR in one simulated individual's V1 at different retinal eccentricities (e.g. where there are no S cones, for foveolal V1, vs. parafoveal). Typically, input from S cones is thought to reduce spatial acuity (because of chromatic aberration). However, such work would just be a step on the way to simulating a more realistic cone mosaic and retinal ganglion cell wiring, as will be discussed next.

Apart from the variable ratios of cones described above, it is also clear (chapter 2) that the arrangement of L and M photoreceptors in human and primate retina is at least partly random, and furthermore they are extremely similar genetically and morphologically. The model of RGC and LGN in this thesis is specifically wired for all cone types, and every cone type exists at each retinal location; that is, cone opponency is perfect. Perfect cone opponency seems implausible to us, but there is debate in the literature about whether or not retinal ganglion cells are indeed wired selectively to L and M cones, which provides a good opportunity to test the effects of random vs. specific wiring in the model.

Initial work has been undertaken in this direction (Kneisel, 2013), and has shown that it will be necessary to consider random wiring's effect on cone opponencies. In central retina, the center pathway of a midget (P) ganglion cell is likely made up of only one cone, while the surround of 5–10, and the surround may contain a mixture of L and M cones. This could change the balance of cone opponencies present. Chapter 5's d_S, the RGC pathway symmetry, indicates that some form of homeostatic mechanism may be required at the RGC and LGN level to ensure a balance of the different opponent channels. Similarly, d_L indicates that the balance of LUM vs. CR signal to the cortex is important for realistic results. Smaller RF sizes will likely favor the chromatic signal, while larger ones the luminance signal (we discuss this further in section 6.4.4 below). Exploring the effects of a realistic cone mosaic and random wiring on these dimensions will be important.

Furthermore, a realistic cone mosaic along with random wiring of the retinal ganglion cells may affect correlations (as mentioned in section 6.4.2 above), though it is not clear how large the effect would be. Instead of L and M cones occupying the same space, as in the models of this thesis, L and M cones would be spatially separate (albeit, very close). Additionally, unlike in our model, the RGC surrounds would not be perfectly overlapping in space. Both of these may reduce correlations between L-M and M-L pathways.

6.4.4 Joint representation of OR and CR

Chapter 5 showed what happens when the balance of OR and CR signal in the model is altered: CR can be organized into patches with OR forming a spatially contiguous map, or CR can be spatially contiguous with OR occupying small regions. This dimension prompts several interesting questions, which we will consider below.

The most obvious unanswered question from our modeling work in chapter 5 is why the particular balance of OR and CR maps? A particular balance has been found in macaque imaging studies, but we do not know for sure if the balance genuinely reflects organization of cells in macaque, or if it comes from experimental procedures. For instance, in the model, a patchy response to RG-BW gratings can be found even when the underlying cells are selective for both CR and OR. However, in the model, we do know that the underlying balance of OR and CR organization can be changed by weighting the subcortical chromatic and achromatic pathways differently, and that we have simply chosen a particular value. Balancing the channels automatically may be possible based on a homeostatic mechanism, but this remains an open question.

Beyond the balance in strength of the two maps as described above, another aspect of their interaction is their spatial scale. In both the naive and idealized models, CR is the dominant map, with OR organized at a smaller scale. The situation in the animal results is not clear, but it would be interesting to investigate the significance of the model's map scales. Miikkulainen et al. (2005) showed in a model of OR and motion direction (DR) selectivity that the relative scales of the OR and DR maps could be altered by changing the speed of the input patterns; at higher speeds, the DR map becomes dominant. The OR and DR model of Miikkulainen et al. (2005) was trained on oriented 2D Gaussian patterns, so the neurons were selective to both orientation and direction. However, in the models of OR and CR in this thesis, an additional complication arises in the interaction of OR and CR: neurons may become strongly selective for CR only, which is another interaction we discuss next.

The subcortical luminance pathway in the model filters out spatially uniform, isoluminant areas of the image, allowing V1 to learn luminance edges. The subcortical chromatic pathway, however, can be activated by spatially uniform, chromatic parts of images (as well as by chromatic edges). Images themselves contain a mixture of luminance edges, chromatic edges, combined luminance and chromatic edges, isoluminant achromatic regions, and isoluminant chromatic regions. The statistics of these in images (Hansen and Gegenfurtner, 2009; Cecchi et al., 2010), combined with the subcortical architecture, will influence how CR and OR are jointly represented in V1. Studying this aspect of the model using both natural images and artificial patterns with controlled statistics, would be interesting for comparison with the range of results in V1 about the separation or otherwise of CR and OR (Shapley and Hawken, 2011).

Finally, in this thesis we have measured orientation preference maps using achromatic gratings, as is typically done in animals. However, (Johnson et al., 2010) measured orientation maps in tree shrew (dichromatic color vision system, with "S" and "ML" cones) using S cone isolating stimuli, finding maps similar to those measured with ML cone isolating stimuli. This would be an interesting result to investigate in the present model, or perhaps more easily in a dichromatic version (see section 6.4.6 below).

6.4.5   Measures of hue map organization

More objective measures of hue map quality may become useful as more data from animals becomes available. Stevens et al. (2013b) computed measures of selectivity and stability for OR maps in a related (achromatic) model, and compared to animal data. It would be possible already to measure selectivity and stability of the model's hue preference over development, but there is not currently any animal data against which to compare. When such data becomes available, it may help to constrain how joint OR and CR maps should develop, as discussed earlier in section 6.2.1.

Another measure that Stevens et al. (2013b) calculate for the related OR model's OR preference maps is pinwheel density. In biological orientation maps across a range of species, the pinwheel density of OR maps has been found to be close to π. Initial work (with Jean-Luc Stevens) indicates that the pinwheel density for hue preference maps of the idealized model presented in chapter 5 is close to 2π (so far, only tested at low values of d_{L}).

6.4.6   Different photoreceptor types

We have modeled the development of perceptual color space in macaque monkey, which, like humans, normally has three distinct cone types. Our present work shows only results from three input channels, but the relative photoreceptor strength dimension could simulate one cone class being absent. This opens up a number of possibilities, which we will discuss below.

Firstly, the resulting model could be used to understand how PCS develops in people who are missing one cone type (which could be the L, M, or S cone), or in regions of cortex where there is no input from S cones (e.g. foveola). Additionally, most mammals are dichromats, including many primates. Modeling dichromatic species may be simpler, and some optical and 2P imaging results are already available (Johnson et al., 2010; Buzás et al., 2008). However, it may be more difficult to relate results from these animals to human perceptual color space, although data from appropriately "color blind" human observers could be used.

In fact, there is a variety of cone sensitivities among humans (chapter 2), both "color normals" and those who are "color blind". It would therefore be very interesting to model the development of PCS using different variations of the standard cone sensitivities, perhaps showing some varieties make little difference, while others are important (e.g. as L and M peaks get closer, the cortex may be less able to distinguish them, and this may affect both the organization of CR and OR in V1). Larger collections of hyperspectral images would help such modeling efforts, because hyperspectral images can easily be used to accurately simulate alternative cone sensitivities.

Previously we mentioned that the addition of a new photopigment type to adult squirrel monkeys led to the monkeys being able to distinguish colors which they previously could not (Mancuso et al., 2009). This experiment could also be simulated, and may provide useful constraints on adaptive mechanisms. In particular, the timescales of activity resulting from expression of the new photopigment vs. observed behavioral changes could be considered.

Finally, it could also be possible to begin modeling rod as well as cone vision. How does the organization of the model cortex change with long-term dim light conditions, or long-term bright light conditions? What effect does a combination of rod and cone input have at intermediate light levels? Also, it has been found that in individuals whose cones are inactive for genetic reasons, the region of cortex normally found to represent solely cones in normal individuals instead represents rods (chapter 2). The region of cortex involved normally represents the foveola, which contains no rods, so

this indicates enough reorganization occurs that a different retinal location is represented (not only a different photoreceptor class). Modeling such phenomena is currently beyond the model's architecture, which makes this suggestion more open than others in this section.

6.4.7 Spatial Adaptation

We saw earlier (section 6.4.1) that one problem caused by the biased hue distribution is for photoreceptor types not to be equally active over the course of a simulation; that problem is solved by homeostasis, as above. The models reviewed in chapter 2 with no adaptation (e.g. the De Paula (2007) model has a fixed one-to-one mapping of hues to photoreceptor activation ratios. One hue presented many more times than another to a model cortex that uses Hebbian learning will cause the cortex to learn that specific ratio of weights more than others, and hence represent that particular hue more than others. Homeostasis of the photoreceptors appears to mitigate this problem to some extent, but not anywhere near as much as the full hue jitter we applied in the idealized model, which causes the full range of different receptor activation ratios to be experienced equally during development. In animals, it is clear that the cones and retinal ganglion cells adapt over multiple timescales, including the short term (chapter 2), meaning there is not a fixed one-to-one relationship between hue and cone activation over the short term. Implementing some form of spatial adaptation in the model may therefore also be important for dealing with the biased distribution of hues experienced during development.

Additionally, a simple model of instantaneous, independent adaptation in each cone class (justified by assuming each cone samples enough locations in a scene, due to eye movements), would make it possible to begin investigating chroma and lightness representation in the model, allowing to move to e.g. CIELCh$_{ab}$ which has a more uniform representation of hue, and which was an unsolved problem from modeling work introduced in chapter 2 (**P6-HSV**). This may also be important for matching conditions of hue map measurement in some experiments (as discussed earlier in section 6.2.1).

6.4.8 Combine with models of other features

We have modeled the development of OR and CR together, emphasizing that it is important to consider more than just CR in isolation. We are most confident about the representation and development over time of OR, but combining only OR and CR

is still a very limited selection of the many features of visual input that appear to be mapped in V1 (chapter 2; Bednar, 2012). To allow a full investigation of some of the suggested psychophysical effects mentioned in this chapter may require extending the subcortical modeling. For instance, some effects transfer between eyes to varying degrees, which can help to constrain where a particular mechanism may be used in the visual pathway (e.g. adaptation of the photoreceptors in one eye cannot explain how an effect could transfer to the other eye). There are also interesting interactions between CR and motion, including apparent motion being induced by color (Favreau et al., 1972), and vice versa. Apart from allowing new adult visual phenomena to be studied, there may also be additional imaging data for other features, which can also help to constrain the model. Additionally, aside from improving the model of CR processing, adding more features also helps to build towards a model of the cortex that can genuinely represent visual input (Gerasymova, 2008; Fisher, 2014).

6.4.9   Model higher cortical areas

The visual pathway does not end at V1, nor is it only a feed-forward pathway. Studies of higher areas of the cortex indicate that various other areas are involved in processing color (see e.g. Conway, 2009), and spatial maps of hue from imaging studies are available in e.g. V2 (Xiao et al., 2003; Lim et al., 2009). Determining whether there is an organization for PCS in areas other than V1 is subject to the same debates as for V1, and furthermore it is not clear how much of the organization for hue in these areas is inherited from previous areas (most importantly, V1). However, a full explanation of PCS will surely involve more than only V1, and the models presented in this thesis will allow future work to begin modeling these areas and their interactions. Even if it eventually turns out that studies finding PCS in V1 need to be reinterpreted in light of new data, the constraints and variabilities shown in this thesis that the brain must be overcoming to develop PCS will remain.

6.5   Conclusion

In this chapter, we first evaluated the work in this thesis against the goals set out in the first chapter. In meeting the goals, we have generated hypotheses about how wiring for the representation of PCS in macaque V1 could develop. We have made suggestions for experimental and modeling work that will help to evaluate our hypotheses, and we

have created two models that serve as a platform from which the modeling work can be conducted.

Chapter 7

Conclusion

This thesis began by making the case that understanding the neural representation of color is particularly useful for understanding how the brain processes information in general, because (a) the difference between perception and physical stimulus is very clear and (b) it is now possible to image the neural activity underpinning perceptual color space at sufficient resolution and spatial scale in primary visual cortex, a relatively early and well understood part of the visual system.

Although experimental work has glimpsed how perceptual color space might be represented in the brain, it has not suggested how the necessary wiring might arise. We review evidence showing that the development of color vision can be strongly affected by both the environment and details of physiology, and yet despite large variabilities and biases in both—between different individuals and for one individual over time—the development of perceptual color space is remarkably consistent and stable.

To understand these findings, we constructed a model of the early visual system with realistic physiology, and using realistic visual input, to see if the wiring could arise through input-driven development. In chapter 4, we showed that the model (which we call the naive model) developed more realistic orientation and color selectivity maps than previous models. However, it did not develop a realistic organization for hue representation. In chapter 5, by analyzing the naive model in detail to learn what prevented it from developing a realistic organization for hue, we showed that a second model, the idealized model, could develop a realistic organization for hue while also maintaining the realistic organization for orientation.

The idealized model clearly shows five dimensions that all previous models of the development of color vision do not take account of: highly non-uniform sampling of perceptual color space experienced during typical natural viewing; variable signal strength

from different cone classes (caused by multiple factors, including highly variable cone ratios, and lens yellowing); highly correlated photoreceptor sensitivities; and how the classical retinal ganglion cell pathway can de-emphasize some hues while also affecting the overall balance of chromatic and achromatic signals. Based on our findings, we hypothesize that even the most basic models of the development of perceptual color space must include several additional adaptive mechanisms in order to be successful, and that such mechanisms are what allows the visual system to be relatively insensitive to many environmental and physiological variabilities and thus robustly develop perceptual color space.

With the naive and idealized models in hand, it is now possible to investigate a wide range of adult color visual phenomena in order to suggest potential adaptive mechanisms, or to further constrain the model. Finally, we have suggested a range of future experimental and modeling work that will help improve the model, and thus further our understanding of how perceptual color space is represented in the brain, along with how that organization develops. This in turn will help to understand both how the adult brain processes information in general, and how it gets that way.

Appendix A

Glossary

achromatic Devoid of hue (CIE, 1987).

brightness Attribute of visual sensation according to which an area appears to emit more or less light (CIE, 1987).

chroma Colorfulness of an area judged as a proportion of the brightness of a similarly illuminated area that appears white or highly transmitting (CIE, 1987).

chromatic Possessing hue (CIE, 1987).

colorfulness Attribute of visual sensation according to which the perceived color appears to be more or less chromatic(CIE, 1987).

cone fundamentals Color-matching functions for human vision (i.e. spectral sensitivity at the cornea). Different from cone photopigment sensitivities, because the fundamentals include filtering effects of the eye's media.

gamut The region enclosed by a color space. The gamut of a color space may be smaller (e.g. sRGB) or larger (e.g. CIELAB) than the gamut perceivable by humans. The human color gamut depends on the cone wavelength sensitivities, so is three dimensional, but gamuts are often presented on two-dimensional chromaticity diagrams such as *(x,y)* diagram by ignoring luminance. Gamut is often used to describe the range of colors a device can display (which depends on its primaries), or to describe the range of colors present in images.

hue Attribute of visual sensation according to which an area appears to be similar to one of the perceived colors: red, yellow, green, and blue, or to a combination of two of them (CIE, 1987).

irradiance Radiant power per unit area incident onto a surface (Wm^{-2}, or $Wm^{-2}nm^{-1}$ if spectral irradiance; Fairchild, 2005).

lightness The brightness of an area judged relative to the brightness of a similarly illuminated area that appears to be white or highly transmitting (CIE, 1987).

luminance Luminance accounts for the energy of the light *and* human spectral sensitivity: $L_v \propto \int P(\lambda)V(\lambda)d\lambda$ where $P(\lambda)$ is radiance and $V(\lambda)$ is the human spectral luminous efficiency function (see figure (c)). The luminance channel in human vision (also referred to as the achromatic channel) is typically taken to be the sum of the L and M photoreceptors.

primaries Three independent light sources, where independent means the color of any one cannot be visually matched by a linear combination of the other two. Primaries may be monochromatic or broadband. A color space requires a set of primaries as part of its definition, and those primaries determine the space's gamut. Note that primaries need not be real (visible); e.g. for CIEXYZ to encompass the complete gamut of colors humans can perceive, its primaries must be more saturated than is possible to realize, since one cone class cannot be stimulated entirely in isolation of another.

radiance See irradiance, but instead of incidence, irradiance measures emission (and consequently has units $\mathrm{Wm^{-2}sr^{-1}}$, or $\mathrm{Wm^{-2}sr^{-1}nm^{-1}}$; (Fairchild, 2005)).

saturation Colorfulness of an area judged in proportion to its brightness (CIE, 1987).

Appendix B

Acronyms

(x,y) diagram CIE 1931 (x, y)-chromaticity diagram.

2AFC two-alternative, forced-choice.

2P in vivo two-photon calcium microscopy.

AL Adaptation, Laterally connected model; Stevens et al., 2013b.

BCID Barcelona Calibrated Image Database; Párraga et al., 2010.

CF connection field.

CIE Commission Internationale de l'Eclairage (International Commission on Illumination).

CIELAB CIE 1976 L* a* b* color space.

CIELCh$_{ab}$ CIE L* C* h$_{ab}$ color space.

CIELUV CIE 1976 L* u* v* color space.

CIEXYZ CIE 1931 XYZ color space.

cmf color-matching functions.

CR color.

DKL Derrington-Krauskopf-Lennie color space.

DoG Difference-of-Gaussians.

DR motion direction.

fMRI functional magnetic resonance imaging.

HSV HSV color space; Smith, 1978.

ICA independent components analysis.

ICID IPL Calibrated Image Database; Laparra et al., 2012.

iOI intrinsic optical imaging.

jnd just-noticeable difference.

LGN lateral geniculate nucleus.

LISSOM Laterally Interconnected Synergetically Self-Organizing Map; Miikkulainen et al., 2005.

LUM luminance.

MCID McGill Calibrated Image Database; Olmos and Kingdom, 2004.

OD ocular dominance.

OI optical imaging.

OR orientation.

PCA principal components analysis.

PCS perceptual color space.

RF receptive field.

RGB RGB color space.

RGC retinal ganglion cells.

SOM self-organizing map; von der Malsburg, 1973.

SPD spectral power distribution.

SP-LMS Smith-Pokorny LMS 2° cone fundamentals.

sRGB sRGB color space; Stokes et al., 1996.

SS-LMS Stockman-Sharpe LMS 2° cone fundamentals.

UCS diagram CIE 1976 uniform chromaticity-scale diagram.

V1 primary visual cortex.

V2 secondary visual cortex.

WS wavelength space.

Appendix C

Parameters

$\mathbf{M_1}$ Matrix to transform CIE 1931 2° XYZ cmf (as modified by Judd 1951 and Vos 1978) into Smith-Pokorny 1975 cone fundamentals: $\begin{pmatrix} 0.2435 & 0.8524 & -0.0516 \\ -0.3954 & 1.1642 & 0.0837 \\ 0 & 0 & 0.6225 \end{pmatrix}$.

$\mathbf{M_2}$ Matrix to transform CIE 1931 2° XYZ cmf (as modified by Judd 1951 and Vos 1978) into sRGB (whitepoint D65): $\begin{pmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & -0.0416 \\ -0.0556 & -0.2040 & 1.0570 \end{pmatrix}$.

m_p Target photoreceptor mean activity: 0.44.

R Maximum circular radius of RGC connection field: 0.375.

γ_ρ Connection strength from CF ρ. For connection between photoreceptor and RGC units, $\gamma_\rho = 4.7$ (or $\gamma_\rho = \frac{4.7}{2}$ where 2 cone types are combined). For V1 afferent, $\gamma_A = 1.7$. For V1 lateral excitatory $\gamma_E = 1.7$, and $\gamma_I = -1.4$ for V1 lateral inhibitory.

m_r Target RGC mean activity: 0.1.

σ_C RGC DoG center Gaussian's radius: 0.036925.

σ_S RGC DoG surround Gaussian's radius: 0.1477.

σ_A See σ_ρ.

σ_ρ Width of Gaussian for V1 projection ρ. For afferent, $\sigma_A = 0.27$. For lateral inhibitory, $\sigma_I = 0.075$. For lateral excitatory, $\sigma_E = 0.025$.

σ_E See σ_ρ.

σ_I See σ_ρ.

α_ρ V1 learning rate for connection type ρ. For afferent, $\alpha_A = 0.1$. For lateral inhibitory, $\alpha_I = 0.3$.

α_A See α_ρ.

α_I See α_ρ.

β Degree of smoothing in calculation of average V1 activity: 0.991.

λ Homeostatic learning rate: 0.01.

μ Target V1 mean activity: 0.024.

r_ρ Maximum circular radius of V1 projection ρ's connection field. For afferent, $r_A = 0.27083$. For lateral inhibitory, $r_I = 0.22917$. For lateral excitatory, $r_E = 0.104$.

r_A See r_ρ.

r_E See r_ρ.

r_I See r_ρ.

Bibliography

Anderson, J. A., and Rosenfeld, E., editors (1988). *Neurocomputing: Foundations of Research*. Cambridge, MA: MIT Press.

Antolik, J. (2010). *Unified Developmental Model of Maps, Complex Cells and Surround Modulation in the Primary Visual Cortex*. Ph.D. thesis, School of Informatics, The University of Edinburgh, UK.

Arbib, M. A., editor (2002). *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA, USA: MIT Press, second edition.

Arieli, A., Grinvald, A., and Slovin, H. (2002). Dural substitute for long-term imaging of cortical activity in behaving monkeys and its clinical implications. *Journal of Neuroscience Methods*, 114(2):119–133.

Atick, J. J., Li, Z., and Redlich, A. N. (1992). Understanding Retinal Color Coding from First Principles. *Neural Computation*, 4(4):559–572.

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.

Baseler, H. A., Brewer, A. A., Sharpe, L. T., Morland, A. B., Jagle, H., and Wandell, B. A. (2002). Reorganization of human cortical maps caused by inherited photoreceptor abnormalities. *Nature Neuroscience*, 5(4):364–370.

Baylor, D. A., Nunn, B. J., and Schnapf, J. L. (1987). Spectral sensitivity of cones of the monkey Macaca fascicularis. *The Journal of Physiology*, 390:145–160.

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

Bednar, J. A., De Paula, J. B., and Miikkulainen, R. (2005). Self-organization of color opponent receptive fields and laterally connected orientation maps. *Neurocomputing*, 65–66:69–76.

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.

Beer, D., Wortman, J., Horwitz, G., and MacLeod, D. (2005). Compensation of white for macular filtering. *Journal of Vision*, 5(8):282.

Belmore, S. C., and Shevell, S. K. (2008). Very-long-term chromatic adaptation: test of gain theory and a new method. *Visual Neuroscience*, 25(3):411–414.

Belmore, S. C., and Shevell, S. K. (2011). Very-long-term and short-term chromatic adaptation: are their influences cumulative? *Vision Research*, 51(3):362–366.

Belpaeme, T., and Bleys, J. (2009). The impact of statistical distributions of colours on colour category acquisition. *Journal of Cognitive Science*, 10(1):1–20.

Benardete, E. A., and Kaplan, E. (1997a). The receptive field of the primate P retinal ganglion cell, II: Nonlinear dynamics. *Visual Neuroscience*, 14(1):187–205.

Benardete, E. A., and Kaplan, E. (1997b). The receptive field of the primate P retinal ganglion cell, I: Linear dynamics. *Visual Neuroscience*, 14(1):169–185.

Benardete, E. A., Kaplan, E., and Knight, B. W. (1992). Contrast gain control in the primate retina: P cells are not X-like, some M cells are. *Visual Neuroscience*, 8:483–486.

Berlin, B., and Kay, P. (1969). *Basic Color Terms: Their Universality and Evolution*. University of California Press.

Blakemore, C., and Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, 228:477–478.

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 Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321(6070):579–585.

Bompas, A., Powell, G., and Sumner, P. (2013). Systematic biases in adult color perception persist despite lifelong information sufficient to calibrate them. *Journal of Vision*, 13(1):19+.

Brainard, D. H. (2003). Color appearance and color difference specification. In Shevell, S. K., editor, *The Science of Color*. Elsevier Science.

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(3):607–614.

Brenner, E., and Cornelissen, F. (2005). A way of selectively degrading colour constancy demonstrates the experience dependence of colour vision. *Current Biology*, 15(21):R864–R866.

Buchsbaum, G., and Gottschalk, A. (1983). Trichromacy, opponent colours coding and optimum colour information transmission in the retina. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 220(1218):89–113.

Burton, G. J., and Moorhead, I. R. (1987). Color and spatial structure in natural scenes. *Applied Optics*, 26(1):157–170.

Buzás, P., Szmajda, B. A., Nezhad, H. M., Dreher, B., and Martin, P. R. (2008). Color signals in the primary visual cortex of marmosets. *Journal of Vision*, 8(10).

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., and Rust, N. C. (2005). Do We Know What the Early Visual System Does? *The Journal of Neuroscience*, 25(46):10577–10597.

Cecchi, G. A., Rao, A. R., Xiao, Y., and Kaplan, E. (2010). Statistics of natural scenes and cortical color processing. *Journal of Vision*, 10(11):21+.

Chakrabarti, A., and Zickler, T. (2011). Statistics of real-world hyperspectral images. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, 193–200. IEEE.

Chapman, B., and Bonhoeffer, T. (1998). Overrepresentation of horizontal and vertical orientation preferences in developing ferret area 17. *Proceedings of the National Academy of Sciences, USA*, 95:2609–2614.

Chapman, B., Stryker, M. P., and Bonhoeffer, T. (1996). Development of orientation preference maps in ferret primary visual cortex. *The Journal of Neuroscience*, 16(20):6443–6453.

Chatterjee, S. (2010). Blobs, micromaps, and two-photon imaging: the functional organization of color in primate V1. In *Society for Neuroscience Abstracts*. Society for Neuroscience, www.sfn.org. Program No. 620.5.

Chatterjee, S., Ohki, K., and Reid, R. C. (2008). Functional micro-architecture of color selectivity in macaque primary visual cortex. In *Society for Neuroscience Abstracts*. Society for Neuroscience, www.sfn.org. Program No. 756.9.

Chichilnisky, E. J., and Wandell, B. A. (1995). Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching. *Vision Research*, 35(2):239–254.

Chklovskii, D. B., and Koulakov, A. A. (2004). Maps in the brain: what can we learn from them? *Annual Review of Neuroscience*, 27:369–392.

CIE (1987). *International Lighting Vocabulary, CIE Publication No. 17.4*. Vienna: CIE.

Ciroux, J. B. (2005). *Simulating the McCollough Effect in a Self-Organizing Model of the Primary Visual Cortex*. Master's thesis, The University of Edinburgh, Scotland, UK.

Clavadetscher, J. E., Brown, A. M., Ankrum, C., and Teller, D. Y. (1988). Spectral sensitivity and chromatic discriminations in 3- and 7-week-old human infants. *Journal of the Optical Society of America A*, 5(12):2093–2105.

Clifford, C. W., Spehar, B., Solomon, S. G., Martin, P. R., and Zaidi, Q. (2003). Interactions between color and luminance in the perception of orientation. *Journal of Vision*, 3(2):106–115.

Conway, B. R. (2003). Colour vision: a clue to hue in V2. *Current Biology*, 13(8):R308–R310.

Conway, B. R. (2009). Color vision, cones, and color-coding in the cortex. *The Neuroscientist*, 15(3):274–290.

Conway, B. R., Chatterjee, S., Field, G. D., Horwitz, G. D., Johnson, E. N., Koida, K., and Mancuso, K. (2010). Advances in Color Science: From Retina to Behavior. *The Journal of Neuroscience*, 30(45):14955–14963.

Conway, B. R., and Livingstone, M. S. (2006). Spatial and temporal properties of cone signals in alert macaque primary visual cortex. *The Journal of Neuroscience*, 26(42):10826–10846.

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.

Crognale, M. A., Nolan, J. B., Webster, M. A., Neitz, M., and Neitz, J. (2001). Color vision and genetics in a case of cone dysfunction syndrome. *Color Research & Application*, 26(S1):284–287.

Crook, J. D., Davenport, C. M., Peterson, B. B., Packer, O. S., Detwiler, P. B., and Dacey, D. M. (2009). Parallel ON and OFF cone bipolar inputs establish spatially coextensive receptive field structure of Blue-Yellow ganglion cells in primate retina. *The Journal of Neuroscience*, 29(26):8372–8387.

Crook, S. M., Davison, A. P., and Plesser, H. E. (2013). Learning from the Past: Approaches for Reproducibility in Computational Neuroscience. In Bower, J. M., editor, *20 Years of Computational Neuroscience*, volume 9 of *Springer Series in Computational Neuroscience*, 73–102. Springer New York.

Curcio, C. A., Allen, K. A., Sloan, K. R., Lerea, C. L., Hurley, J. B., Klock, I. B., and Milam, A. H. (1991). Distribution and morphology of human cone photoreceptors stained with anti-blue opsin. *The Journal of Comparative Neurology*, 312(4):610–624.

Cynader, M., and Chernenko, G. (1976). Abolition of direction selectivity in the visual cortex of the cat. *Science*, 193(4252):504–505.

Dacey, D. M. (2000). Parallel pathways for spectral coding in primate retina. *Annual Review of Neuroscience*, 23:743–775.

Dan, Y., Atick, J. J., and Reid, R. C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *The Journal of Neuroscience*, 16(10):3351–3362.

De Paula, J. B. (2007). *Modeling the Self-Organization of Color Selectivity in the Visual Cortex*. Ph.D. thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX.

De Valois, R. L., and De Valois, K. K. (1993). A multi-stage color model. *Vision Research*, 33(8):1053–1065.

De Valois, R. L., Morgan, H. C., Polson, M. C., Mead, W. R., and Hull, E. M. (1974). Psychophysical studies of monkey vision—I. Macaque luminosity and color vision tests. *Vision Research*, 14(1):53–67.

Delahunt, P. B., Webster, M. A., Ma, L., and Werner, J. S. (2004). Long-term renormalization of chromatic mechanisms following cataract surgery. *Visual Neuroscience*, 21(3):301–307.

Derrington, A. M., Krauskopf, J., and Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *The Journal of Physiology*, 357(1):241–265.

Doi, E. (2003). *A Study of Computational Neural Network Models on Spatio-Chromatic Properties of the Early Visual System*. Ph.D. thesis, Kyoto University, Kyoto, Japan.

Doi, E., Inui, T., Won, T., Wachtler, T., and Sejnowski, T. J. (2003). Spatiochromatic receptive field properties derived from information-theoretic analyses of cone mosaic responses to natural scenes. *Neural Computation*, 15(2):397–417.

Dubois, P. F., Hinsen, K., and Hugunin, J. (1996). Numerical Python. *Computers in Physics*, 10(3):262–267.

Dunn, F. A., Lankheet, M. J., and Rieke, F. (2007). Light adaptation in cone vision involves switching between receptor and post-receptor sites. *Nature*, 449(7162):603–606.

Eguchi, A., Neymotin, S. A., and Stringer, S. M. (2014). Color opponent receptive fields self-organize in a biophysical model of visual cortex via spike-timing dependent plasticity. *Frontiers in Neural Circuits*, 8(16).

Eisner, A., and Enoch, J. (1982). Some effects of 1 week's monocular exposure to long-wavelength stimuli. *Perception & Psychophysics*, 31(2):169–174.

Ellis, S. R. (1977). Orientation selectivity of the McCollough effect: Analysis by equivalent contrast transformation. *Perception and Psychophysics*, 22(6):539–544.

Endeman, D., and Kamermans, M. (2010). Cones perform a non-linear transformation on natural stimuli. *The Journal of Physiology*, 588(Pt 3):435–446.

Fairchild, M. D. (2005). *Color Appearance Models*. Wiley-Interscience.

Fairman, H. S., Brill, M. H., and Hemmendinger, H. (1997). How the CIE 1931 color-matching functions were derived from Wright-Guild data. *Color Research & Application*, 22(1):11–23.

Favreau, O. E., Emerson, V. F., and Corballis, M. C. (1972). Motion perception: a color-contingent aftereffect. *Science*, 176(4030):78–79.

Fawcett, S. L., Wang, Y.-Z. Z., and Birch, E. E. (2005). The critical period for susceptibility of human stereopsis. *Investigative Ophthalmology & Visual Science*, 46(2):521–525.

Field, G. D., Gauthier, J. L., Sher, A., Greschner, M., Machado, T. A., Jepson, L. H., Shlens, J., Gunning, D. E., Mathieson, K., Dabrowski, W., Paninski, L., Litke,

A. M., and Chichilnisky, E. J. (2010). Functional connectivity in the retina at the resolution of photoreceptors. *Nature*, 467(7316):673–677.

Field, G. D., Sher, A., Gauthier, J. L., Greschner, M., Shlens, J., Litke, A. M., and Chichilnisky, E. J. (2007). Spatial properties and functional organization of small bistratified ganglion cells in primate retina. *The Journal of Neuroscience*, 27(48):13261–13272.

Finlayson, G. D., Drew, M. S., and Funt, B. V. (1994). Spectral sharpening: sensor transformations for improved color constancy. *Journal of the Optical Society of America A*, 11(5):1553–1563.

Fisher, T. (2014). *Model of all known spatial maps in primary visual cortex*. Master's thesis, The University of Edinburgh, Scotland, UK.

Foster, D. H. (2010). Chromatic Function of the Cones. In Dana, R., and Bex, P., editors, *Encyclopedia of the Eye*, 266–274. Academic Press, first edition.

Foster, D. H., Amano, K., Nascimento, S. M., and Foster, M. J. (2006). Frequency of metamerism in natural scenes. *Journal of the Optical Society of America A*, 23(10):2359–2372.

Friedman, H. S., Zhou, H., and von der Heydt, R. (2003). The coding of uniform colour figures in monkey visual cortex. *The Journal of Physiology*, 548(Pt 2):593–613.

Garrigan, P., Ratliff, C. P., Klein, J. M., Sterling, P., Brainard, D. H., and Balasubramanian, V. (2010). Design of a trichromatic cone array. *PLoS Computational Biology*, 6(2):e1000677+.

Gegenfurtner, K. (2001). Color in the cortex revisited. *Nature Neuroscience*, 4(4):339–340.

Gegenfurtner, K. R. (2003). Cortical mechanisms of colour vision. *Nature Reviews Neuroscience*, 4(7):563–572.

Gegenfurtner, K. R., and Kiper, D. C. (2003). Color vision. *Annual Review of Neuroscience*, 26:181–206.

Geisler, W. S., and Perry, J. S. (2011). Statistics for optimal point prediction in natural images. *Journal of Vision*, 11(12):14+.

Gerasymova, K. (2008). *Computational Model of All Known Feature Maps in the Primary Visual Cortex*. Master's thesis, Humboldt-Universitaet zu Berlin, Germany.

Gilbert, C. D., Li, W., and Piech, V. (2009). Perceptual learning and adult cortical plasticity. *The Journal of Physiology*, 587(12):2743–2751.

Goddard, E., Mannion, D. J., McDonald, J. S., Solomon, S. G., and Clifford, C. W. G. (2010). Combination of subcortical color channels in human visual cortex. *Journal of Vision*, 10(5):25+.

Granzier, J. J. M., and Valsecchi, M. (2014). Variations in daylight as a contextual cue for estimating season, time of day, and weather conditions. *Journal of Vision*, 14(1):22+.

Hamer, R. D., Alexander, K. R., and Teller, D. Y. (1982). Rayleigh discriminations in young human infants. *Vision Research*, 22(5):575–577.

Hansen, T., and Gegenfurtner, K. R. (2009). Independence of color and luminance edges in natural scenes. *Visual Neuroscience*, 26(1):35–49.

Hendley, C. D., and Hecht, S. (1949). The colors of natural objects and terrains, and their relation to visual color deficiency. *Journal of the Optical Society of America*, 39(10):870–873.

Hofer, H., Carroll, J., Neitz, J., Neitz, M., and Williams, D. R. (2005a). Organization of the human trichromatic cone mosaic. *The Journal of Neuroscience*, 25(42):9669–9679.

Hofer, H., Singer, B., and Williams, D. R. (2005b). Different sensations from cones with the same photopigment. *Journal of Vision*, 5(5):5–454.

Horton, J. C., and Adams, D. L. (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456):837–862.

Hubel, D. H., and Wiesel, T. N. (1959). Receptive fields of single neurons in the cat's striate cortex. *The Journal of Physiology*, 148:574–591.

Hubel, D. H., and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, 195:215–243.

Huberman, A. D., Feller, M. B., and Chapman, B. (2008). Mechanisms underlying development of visual maps and receptive fields. *Annual Review of Neuroscience*, 31:479–509.

Hurvich, L. M., and Jameson, D. (1957). An opponent-process theory of color vision. *Psychological review*, 64, Part 1(6):384–404.

Jacobs, G. H. (2008). Primate color vision: a comparative perspective. *Visual Neuroscience*, 25(5-6):619–633.

Jacobs, G. H. (2012). The evolution of vertebrate color vision. *Advances in Experimental Medicine and Biology*, 739:156–172.

Jameson, D., Hurvich, L. M., and Varner, F. D. (1979). Receptoral and postreceptoral visual processes in recovery from chromatic adaptation. *Proceedings of the National Academy of Sciences*, 76(6):3034–3038.

Johnson, E. N., Hawken, M. J., and Shapley, R. (2001). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience*, 4(4):409–416.

Johnson, E. N., Hawken, M. J., and Shapley, R. (2008). The orientation selectivity of color-responsive neurons in macaque V1. *The Journal of Neuroscience*, 28(32):8096–8106.

Johnson, E. N., Van Hooser, S. D., and Fitzpatrick, D. (2010). The representation of S-cone signals in primary visual cortex. *The Journal of Neuroscience*, 30(31):10337–10350.

Kaplan, E., and Benardete, E. (2001). The dynamics of primate retinal ganglion cells. *Progress in Brain Research*, 134:17–34.

Kitaoka, A. (2008). Red and yellow beads. `www.psy.ritsumei.ac.jp/˜akitaoka/zanzo2e.html`. Accessed: 2014-08-17.

Kitaoka, A. (2012). Ritsumeikan flags. `www.psy.ritsumei.ac.jp/˜akitaoka/colorconstancy2e.html`. Accessed: 2014-08-17.

Kitaoka, A. (2013). Eye color constancy signal. `www.psy.ritsumei.ac.jp/˜akitaoka/colorconstancy2e.html`. Accessed: 2014-08-17.

Kneisel, Y. (2013). *Feasibility of Random Wiring in Primate Colour Vision*. Master's thesis, The University of Edinburgh, UK.

Krantz, D. H. (1975). Color measurement and color theory: I. Representation theorem for Grassmann structures. *Journal of Mathematical Psychology*, 12(3):283–303.

Kröger, R. H., Braun, S. C., and Wagner, H. (2001). Rearing in different photic and chromatic environments modifies spectral responses of cone horizontal cells in adult fish retina. *Visual Neuroscience*, 18(6):857–864.

Kröger, R. H. H., Bowmaker, J. K., and Wagner, H. J. (1999). Morphological changes in the retina of Aequidens pulcher (Cichlidae) after rearing in monochromatic light. *Vision Research*, 39(15):2441–2448.

Kuchenbecker, J. A., Sahay, M., Tait, D. M., Neitz, M., and Neitz, J. (2008). Topography of the long- to middle-wavelength sensitive cone ratio in the human retina assessed with a wide-field color multifocal electroretinogram. *Visual Neuroscience*, 25(3):301–306.

Kuehni, R. G. (2004). Variability in unique hue selection: a surprising phenomenon. *Color Research & Application*, 29(2):158–162.

Kuehni, R. G. (2014). Unique hues and their stimuli—state of the art. *Color Research & Application*, 39(3):279–287.

Lafer-Sousa, R., Liu, Y. O., Lafer-Sousa, L., Wiest, M. C., and Conway, B. R. (2012). Color tuning in alert macaque V1 assessed with fMRI and single-unit recording shows a bias toward daylight colors. *Journal of the Optical Society of America A*, 29(5):657–670.

Landisman, C. E., and Ts'o, D. Y. (2002). Color processing in macaque striate cortex: Relationships to ocular dominance, cytochrome oxidase, and orientation. *Journal of Neurophysiology*, 87(6):3126–3137.

Laparra, V., Jiménez, S., Valls, G. C., and Malo, J. (2012). Nonlinearities and adaptation of color vision from sequential principal curves analysis. *Neural Computation*, 24(10):2751–2788.

Lee, B. B., Pokorny, J., Smith, V. C., Martin, P. R., and Valberg, A. (1990). Luminance and chromatic modulation sensitivity of macaque ganglion cells and human observers. *Journal of the Optical Society of America A*, 7(12):2223–2236.

Lee, B. B., Smith, V. C., Pokorny, J., and Sun, H. (2008). Chromatic adaptation in red-green cone-opponent retinal ganglion cells of the macaque. *Vision Research*, 48(26):2625–2632.

Lee, R. J., Dawson, K. A., and Smithson, H. E. (2012). Slow updating of the achromatic point after a change in illumination. *Journal of Vision*, 12(1):19+.

Lee, R. L., and Hernández-Andrés, J. (2005). Colors of the daytime overcast sky. *Applied Optics*, 44(27):5712–5722.

Lennie, P., Haake, P. W., and Williams, D. R. (1991). The Design of Chromatically Opponent Receptive Fields. In Landy, M. S., and Movshon, J. A., editors, *Computational Models of Visual Processing*, 71–82. Cambridge, MA: MIT Press.

Lennie, P., Krauskopf, J., and Sclar, G. (1990). Chromatic mechanisms in striate cortex of macaque. *The Journal of Neuroscience*, 10(2):649–669.

Lewis, A., and Zhaoping, L. (2006). Are cone sensitivities determined by natural color statistics? *Journal of Vision*, 6(3):8–302.

Li, M., Liu, F., Juusola, M., and Tang, S. (2014). Perceptual color map in macaque visual area V4. *The Journal of Neuroscience*, 34(1):202–217.

Li, Y., Fitzpatrick, D., and White, L. E. (2006). The development of direction selectivity in ferret visual cortex requires early visual experience. *Nature Neuroscience*, 9(5):676–681.

Li, Y., Hooser, S. D. V., Mazurek, M., White, L. E., and Fitzpatrick, D. (2008). Experience with moving visual stimuli drives the early development of cortical direction selectivity. *Nature*, 456(7224):952–956.

Lim, H., Wang, Y., Xiao, Y., Hu, M., and Felleman, D. J. (2009). Organization of hue selectivity in macaque V2 thin stripes. *Journal of Neurophysiology*, 102(5):2603–2615.

Linhares, J. a. M., Pinto, P. D., and Nascimento, S. M. (2008). The number of discernible colors in natural scenes. *Journal of the Optical Society of America A*, 25(12):2918–2924.

Livingstone, M., and Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, 240(4853):740–749.

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.

Logvinenko, A. D. (2012). A theory of unique hues and colour categories in the human colour vision. *Color Research & Application*, 37(2):109–116.

Lotto, R., and Purves, D. (2002). A rationale for the structure of color space. *Trends in Neurosciences*, 25(2):84–88.

Lovell, P. G., Tolhurst, D. J., Párraga, C. A., Baddeley, R., Leonards, U., Troscianko, J., and Troscianko, T. (2005). Stability of the color-opponent signals under changes of illuminant in natural scenes. *Journal of the Optical Society of America A*, 22(10):2060–2071.

Lu, H. D., and Roe, A. W. (2008). Functional organization of color domains in V1 and V2 of macaque monkey revealed by optical imaging. *Cerebral Cortex*, 18(3):516–533.

MacEvoy, B. (2005). Modern color models. `http://www.handprint.com/HP/WCL/color7.html`. Accessed: 2014-08-01.

MacEvoy, B. (2009). Light and the eye. `http://www.handprint.com/HP/WCL/color1.html`. Accessed: 2014-08-01.

Magnussen, S., Spillmann, L., Stürzel, F., and Werner, J. S. (2001). Filling-in of the foveal blue scotoma. *Vision Research*, 41(23):2961–2967.

Mancuso, K., Hauswirth, W. W., Li, Q., Connor, T. B., Kuchenbecker, J. A., Mauck, M. C., Neitz, J., and Neitz, M. (2009). Gene therapy for red-green colour blindness in adult primates. *Nature*, 461(7265):784–787.

Mancuso, K., Mauck, M. C., Kuchenbecker, J. A., Neitz, M., and Neitz, J. (2010). A multi-stage color model revisited: implications for a gene therapy cure for red-green colorblindness. *Advances in Experimental Medicine and Biology*, 664:631–638.

Marín-Franch, I., and Foster, D. H. (2010). Number of perceptually distinct surface colors in natural scenes. *Journal of Vision*, 10(9):9+.

Masaoka, K., Berns, R. S., Fairchild, M. D., and Moghareh Abed, F. (2013). Number of discernible object colors is a conundrum. *Journal of the Optical Society of America A*, 30(2):264–277.

McCollough Howard, C., and Webster, M. A. (2011). McCollough effect. *Scholarpedia*, 6(2):8175. revision #127566.

McCourt, M. E., and Jacobs, G. H. (1983). Effects of photic environment on the development of spectral response properties of optic nerve fibers in the ground squirrel. *Experimental Brain Research*, 49(3):443–452.

McDermott, K. C., Malkoc, G., Mulligan, J. B., and Webster, M. A. (2010). Adaptation and visual salience. *Journal of Vision*, 10(13).

McDermott, K. C., and Webster, M. A. (2012). Uniform color spaces and natural image statistics. *Journal of the Optical Society of America A*, 29(2).

McLelland, D., Ahmed, B., and Bair, W. (2009). Responses to static visual images in macaque lateral geniculate nucleus: implications for adaptation, negative afterimages, and visual fading. *The Journal of Neuroscience*, 29(28):8996–9001.

Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (2005). *Computational Maps in the Visual Cortex*. Berlin: Springer.

Miyahara, E., Pokorny, J., Smith, V. C., Baron, R., and Baron, E. (1998). Color vision in two observers with highly biased LWS/MWS cone ratios. *Vision Research*, 38(4):601–612.

Mizokami, Y., Werner, J. S., Crognale, M. A., and Webster, M. A. (2006). Nonlinearities in color coding: compensating color appearance for the eye's spectral sensitivity. *Journal of Vision*, 6(9):996–1007.

Naka, K. I., and Rushton, W. A. (1966). S-potentials from luminosity units in the retina of fish (Cyprinidae). *The Journal of Physiology*, 185(3):587–599.

Nascimento, S. M. C., Linhares, J. M. M., Pinto, P. D., Foster, D. H., and Amano, K. (2008). The distribution of discernible colours in natural scenes. In *ECVP Abstract Supplement*, volume 37, 84. Pion Ltd.

Nassau, K. (2003). The physics and chemistry of color: the 15 mechanisms. In Shevell, S. K., editor, *The Science of Color*. Elsevier Science.

Neitz, J., Carroll, J., Yamauchi, Y., Neitz, M., and Williams, D. R. (2002). Color perception is mediated by a plastic neural mechanism that is adjustable in adults. *Neuron*, 35(4):783–792.

Neitz, J., and Neitz, M. (2008). Colour vision: the wonder of hue. *Current Biology*, 18(16):R700–R702.

Neitz, J., and Neitz, M. (2011). The genetics of normal and defective color vision. *Vision Research*, 51(7):633–651.

Nickerson, D., Kelly, K. L., and Stultz, K. F. (1945). Color of soils. *Journal of the Optical Society of America*, 35(4):297.

Ohki, K., Chatterjee, S., and Reid, R. C. (2008). Functional micro-architecture of orientation selectivity in macaque primary visual cortex. In *Society for Neuroscience Abstracts*. Society for Neuroscience, www.sfn.org. Program No. 756.9.

Ohki, K., Chung, S., Kara, P., Hubener, M., Bonhoeffer, T., and Reid, R. C. (2006). Highly ordered arrangement of single neurons in orientation pinwheels. *Nature*, 442(7105):925–928.

Ohki, K., and Reid, R. C. (2007). Specificity and randomness in the visual cortex. *Current Opinion in Neurobiology*, 17(4):401–407.

Oliphant, T. E. (2007). Python for Scientific Computing. *Computing in Science and Engineering*, 9(3):10–20.

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. (2005). How Close Are We to Understanding V1? *Neural Computation*, 17(8):1665–1699.

O'Neil, S. F., and Webster, M. A. (2014). Filling in, filling out, or filtering out: processes stabilizing color appearance near the center of gaze. *Journal of the Optical Society of America A*, 31(4):A140–A147.

Palmer, C. M. (2009). *Topographic and Laminar Models for the Development and Organisation of Spatial Frequency and Orientation in V1*. Ph.D. thesis, School of Informatics, The University of Edinburgh, UK.

Párraga, C. A., Baldrich, R., and Vanrell, M. (2010). Accurate mapping of natural scenes radiance to cone activation space: a new image dataset. *Conference on Colour in Graphics, Imaging, and Vision*, 50–57.

Párraga, C. A., Brelstaff, G., Troscianko, T., and Moorehead, I. R. (1998). Color and luminance information in natural scenes. *Journal of the Optical Society of America A*, 15(3):563–569.

Parry, N. R. A., McKeefry, D. J., and Murray, I. J. (2006). Variant and invariant color perception in the near peripheral retina. *Journal of the Optical Society of America A*, 23(7):1586–1597.

Pitkow, X., and Meister, M. (2012). Decorrelation and efficient coding by retinal ganglion cells. *Nature Neuroscience*, 15(4):628–635.

Pokorny, J., and Smith, V. (2013). Spectral sensitivities of the human cones. *Journal of Vision*, 13(15):T20.

Preuss, T. (2004). Specializations of the human visual system: The monkey model meets human reality. In Kaas, J., and Collins, C., editors, *The Primate Visual System*, 231–259. Florida: CRC Press.

Purdy, D. M. (1931). Spectral hue as a function of intensity. *American Journal of Psychology*, 43(4):pp. 541–559.

Rao, A. R., Cecchi, G., Peck, C., and Kozloski, J. (2005). A model of the formation of a self-organized cortical representation of color. In Rogowitz, B. E., Pappas, T. N., and Daly, S. J., editors, *Human Vision and Electronic Imaging X*, volume 5666, 17–26. SPIE.

Rao, A. R., and Xiao, Y. (2012). A computational model of early visual cortex using konio-cellular pathway projections. In *Neural Networks (IJCNN), The 2012 International Joint Conference on*, 1–8. IEEE.

Regan, B. C., Julliot, C., Simmen, B., Viénot, F., Charles-Dominique, P., and Mollon, J. D. (2001). Fruits, foliage and the evolution of primate colour vision. *Philosophical Transactions of the Royal Society B*, 356(1407):229–283.

Rieke, F., and Rudd, M. E. (2009). The challenges natural images pose for visual adaptation. *Neuron*, 64(5):605–616.

Roorda, A., Metha, A. B., Lennie, P., and Williams, D. R. (2001). Packing arrangement of the three cone classes in primate retina. *Vision Research*, 41(10-11):1291–1306.

Roorda, A., and Williams, D. R. (1999). The arrangement of the three cone classes in the living human eye. *Nature*, 397(6719):520–522.

Ruderman, D. L., Cronin, T. W., and Chiao, C.-C. (1998). Statistics of cone responses

to natural images: implications for visual coding. *Journal of the Optical Society of America A*, 15(8):2036–2045.

Saarinen, J., and Kohonen, T. (1985). Self-organized formation of colour maps in a model cortex. *Perception*, 14:711–719.

Schefrin, B. E., Adams, A. J., and Werner, J. S. (1991). Anomalies beyond sites of chromatic opponency contribute to sensitivity losses of an S-cone pathway in diabetes. *Clinical Vision Sciences*, 6:219–228.

Schefrin, B. E., and Werner, J. S. (1990). Loci of spectral unique hues throughout the life span. *Journal of the Optical Society of America A*, 7(2):305–311.

Sclar, G., and Freeman, R. D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research*, 46(3):457–461.

Shapley, R., and Hawken, M. J. (2011). Color in the cortex: single- and double-opponent cells. *Vision Research*, 51(7):701–717.

Shepard, R. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. II. *Psychometrika*, 27(3):219–246.

Shevell, S. K. (2003). Color appearance. In Shevell, S. K., editor, *The Science of Color*. Elsevier Science.

Shouval, H. (2007). Models of synaptic plasticity. Scholarpedia, 2(7):1605. revision #91520.

Smith, A. R. (1978). Color Gamut Transform Pairs. In *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques*, volume 12 of *SIGGRAPH '78*, 12–19. New York, NY, USA: ACM.

Smith, T., and Guild, J. (1931). The C.I.E. colorimetric standards and their use. *Transactions of the Optical Society*, 33(3):73+.

Smith, V. C., and Pokorny, J. (1975). Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm. *Vision Research*, 15(2):161–171.

Solomon, S. G., and Lennie, P. (2005). Chromatic gain controls in visual cortical neurons. *The Journal of Neuroscience*, 25(19):4779–4792.

Solomon, S. G., and Lennie, P. (2007). The machinery of colour vision. *Nature Reviews Neuroscience*, 8(4):276–286.

Solomon, S. G., Tailby, C., Cheong, S. K., and Camp, A. J. (2010). Linear and non-linear contributions to the visual sensitivity of neurons in primate lateral geniculate nucleus. *Journal of Neurophysiology*, 104(4):1884–1898.

Spigler, G. (2014). *Neural Modelling of the McCollough Effect in colour vision*. Master's thesis, The University of Edinburgh, Scotland, UK.

Stanikunas, R., Vaitkevicius, H., and Kulikowski, J. J. (2004). Investigation of color constancy with a neural network. *Neural Networks*, 17:327–337.

Stevens, J.-L. R., Elver, M. I., and Bednar, J. A. (2013a). An automated and reproducible workflow for running and analyzing neural simulations using Lancet and IPython Notebook. *Frontiers in Neuroinformatics*, 7:44.

Stevens, J.-L. R., Law, J. S., Antolik, J., and Bednar, J. A. (2013b). Mechanisms for stable, robust, and adaptive development of orientation maps in the primary visual cortex. *The Journal of Neuroscience*, 33:15747–15766.

Stockman, A., and Sharpe, L. (1999). Cone spectral sensitivities and color matching. In Gegenfurtner, K., and Sharpe, L., editors, *Color vision: from Genes to Perception*, 53–87. Cambridge: Cambridge University Press.

Stokes, M., Anderson, M., Chandrasekar, S., and Motta, R. (1996). A standard default color space for the Internet—sRGB. `www.w3.org/Graphics/Color/sRGB`. Accessed: 2014-08-17.

Stoughton, C. M., Lafer-Sousa, R., Gagin, G., and Conway, B. R. (2012). Psychophysical chromatic mechanisms in macaque monkey. *The Journal of Neuroscience*, 32(43):15216–15226.

Strangor, C. (2010). *Introduction to Psychology*. Flat World Knowledge.

Stromeyer, C. F. (1969). Further studies of the McCollough effect. *Perception & Psychophysics*, 6(2):105–110.

Sugita, Y. (2004). Experience in early infancy is indispensable for color perception. *Current Biology*, 14(14):1267–1271.

Swindale, N. (2008). Visual map. Scholarpedia, 3(6):4607. revision #91925.

Swindale, N. V. (2004). How different feature spaces may be represented in cortical maps. *Network: Computation in Neural Systems*, 15(4):217–242.

Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., and Hubener, M. (2000). Visual cortex maps are optimized for uniform coverage. *Nature Neuroscience*, 3(8):822–826.

Takamura, S., and Kobayashi, N. (2001). Constructing a uniform color space for visually lossless color representation and image coding. In *Image Processing, 2001. Proceedings. 2001 International Conference on*, volume 1, 918–921 vol.1. IEEE.

Tavazoie, S. F., and Reid, R. C. (2000). Diverse receptive fields in the lateral geniculate nucleus during thalamocortical development. *Nature Neuroscience*, 3(6):608–616.

Teller, D. Y. (1998). Spatial and temporal aspects of infant color vision. *Vision Research*, 38(21):3275–3282.

Thompson, P., and Latchford, G. (1986). Colour-contingent after-effects are really wavelength-contingent. *Nature*, 320(6062):525–526.

Tkačik, G., Garrigan, P., Ratliff, C., Milčinski, G., Klein, J. M., Seyfarth, L. H., Sterling, P., Brainard, D. H., and Balasubramanian, V. (2011). Natural images from the birthplace of the human eye. *PLoS ONE*, 6(6).

Valeton, J. M., and van Norren, D. (1983). Light adaptation of primate cones: An analysis based on extracellular data. *Vision Research*, 23(12):1539–1547.

Vazquez-Corral, J., and Bertalmío, M. (2014). Spectral sharpening of color sensors: diagonal color constancy and beyond. *Sensors (Basel, Switzerland)*, 14(3):3965–3985.

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.

Wachtler, T., Lee, T.-W., and Sejnowski, T. J. (2001). Chromatic structure of natural scenes. *Journal of the Optical Society of America A*, 18(1):65–77.

Wachtler, T., Sejnowski, T. J., and Albright, T. D. (2003). Representation of color stimuli in awake macaque primary visual cortex. *Neuron*, 37(4):681–691.

Wagner, H.-J., and Kröger, R. H. H. (2000). Effects of long-term spectral deprivation on the morphological organization of the outer retina of the blue acara (Aequidens pulcher). *Philosophical Transactions of the Royal Society B*, 355(1401):1249–1252.

Wagner, H.-J., and Kröger, R. H. H. (2005). Adaptive plasticity during the development of colour vision. *Progress in Retinal and Eye Research*, 24(4):521–536.

Walraven, J., and Werner, J. S. (1991). The invariance of unique white; a possible implication for normalizing cone action spectra. *Vision Research*, 31(12):2185–2193.

Wandell, B. A. (1995). *Foundations of Vision*. Sinauer Associates Inc.

Wark, B., Lundstrom, B. N., and Fairhall, A. (2007). Sensory adaptation. *Current Opinion in Neurobiology*, 17(4):423–429.

Webster, M., Miyahara, E., Malkoc, G., and Raker, V. (2000). Variations in normal color vision. II. unique hues. *Journal of the Optical Society of America A*, 17(9):1545–1555.

Webster, M., and Mollon, J. (1991). Changes in colour appearance following post-receptoral adaptation. *Nature*, 349(0):235–238.

Webster, M. A. (2011). Adaptation and visual coding. *Journal of Vision*, 11(5).

Webster, M. A., Halen, K., Meyers, A. J., Winkler, P., and Werner, J. S. (2010). Colour appearance and compensation in the near periphery. *Proceedings of the Royal Society B*, 277(1689):1817–1825.

Webster, M. A., Mizokami, Y., and Webster, S. M. (2007). Seasonal variations in the color statistics of natural images. *Network: Computation in Neural Systems*, 18(3):213–233.

Webster, M. A., and Mollon, J. D. (1997). Adaptation and the color statistics of natural images. *Vision Research*, 37(23):3283–3298.

Webster, M. A., Webster, S. M., Bharadwaj, S., Verma, R., Jaikumar, J., Madan, G., and Vaithilingham, E. (2002). Variations in normal color vision. III. unique hues in indian and united states observers. *Journal of the Optical Society of America A*, 19(10):1951–1962.

Werner, J. S., Delahunt, P. B., and Hardy, J. L. (2004). Chromatic-Spatial vision of the aging eye. *Optical Review*, 11(4):226–234.

Werner, J. S., and Schefrin, B. E. (1993). Loci of achromatic points throughout the life span. *Journal of the Optical Society of America A*, 10(7):1509–1516.

Westheimer, G. (2007). The ON-OFF dichotomy in visual processing: from receptors to perception. *Progress in Retinal and Eye Research*, 26(6):636–648.

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

Williams, D. R., MacLeod, D. I., and Hayhoe, M. M. (1981). Foveal tritanopia. *Vision Research*, 21(9):1341–1356.

Wilson, S. P., and Bednar, J. A. (2015). What, if anything, are topological maps for? *Developmental Neurobiology*. In press.

Wuerger, S. (2013). Colour constancy across the life span: evidence for compensatory mechanisms. *PLoS ONE*, 8(5).

Wuerger, S. M., Maloney, L. T., and Krauskopf, J. (1995). Proximity judgments in color space: tests of a Euclidean color geometry. *Vision Research*, 35(6):827–835.

Wyszecki, G., and Stiles, W. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley-Interscience.

Xiao, Y., Casti, A., Xiao, J., and Kaplan, E. (2007). Hue maps in primate striate cortex. *Neuroimage*, 35(2):771–786.

Xiao, Y., Wang, Y., and Felleman, D. J. (2003). A spatially organized representation of color in macaque cortical area V2. *Nature*, 421:535–539.

Yeh, T., Lee, B. B., and Kremers, J. (1996). The time course of adaptation in macaque retinal ganglion cells. *Vision Research*, 36(7):913–931.

Yendrikhovskij, S. N. (2001a). A computational model of colour categorization. *Color Research & Application*, 26:S235–S238.

Yendrikhovskij, S. N. (2001b). Computing color categories from statistics of natural images. *Journal of Imaging Science and Technology*, 409–417.

Yu, H., Farley, B. J., Jin, D. Z., and Sur, M. (2005). The coordinated mapping of visual space and response features in visual cortex. *Neuron*, 47(2):267–280.

Zaidi, Q. (1997). Decorrelation of L- and M-cone signals. *Journal of the Optical Society of America A*, 14(12):3430–3431.

Zaidi, Q., Victor, J., McDermott, J., Geffen, M., Bensmaia, S., and Cleland, T. A. (2013). Perceptual spaces: mathematical structures to neural mechanisms. *The Journal of Neuroscience*, 33(45):17597–17602.

Zeki, S. (1983). Colour coding in the cerebral cortex: the reaction of cells in monkey visual cortex to wavelengths and colours. *Neuroscience*, 9(4):741–765.

Zhang, X., and Wandell, B. A. (1997). A spatial extension of CIELAB for digital color-image reproduction. *Journal of the Society for Information Display*, 5(1):61–63.