UNIVERSITY OF CALIFORNIA,
IRVINE

Psychophysics of Color Vision

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Psychology

by

Christian Herrera

Dissertation Committee:
Professor Charles F. Chubb, Chair
Professor George Sperling
Professor Charles E. Wright

2016

DEDICATION

To my parents, Agustin and Pilar
To my partner, Erika
To my son, Matias

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank Charlie, for his indefatigable support, and for exemplary scholarship and mentoring.

I would like to thank Ted and George, for exemplary scholarship.

I would like to thank Consejo Nacional de Ciencia y Tecnologia and UC-Mexus, for a five year grant.

I would like to thank the Department of Cognitive Sciences at UC Irvine.

CURRICULUM VITAE

Christian Herrera

EDUCATION

Doctor of Philosophy in Psychology **2016**
University of California, Irvine *Irvine, California*

Master of Science in Cognitive Neuroscience **2013**
University of California, Irvine *Irvine, California*

Bachelor of Arts in Psychology **2009**
University of California, Irvine *Irvine, California*

RESEARCH EXPERIENCE

Graduate Research Assistant **2007–2012**
University of California, Irvine *Irvine, California*

TEACHING EXPERIENCE

Teaching Assistant **2010–2013**
University of California, Irvine *Irvine, California*

ABSTRACT OF THE DISSERTATION

Psychophysics of Color Vision

By

Christian Herrera

Doctor of Philosophy in Psychology

University of California, Irvine, 2016

Professor Charles F. Chubb, Chair

All visual information available to us is the result of combining signals from photoreceptors in the retina differentially sensitive to three wavelengths. Such information processing occurs at different stages, and current models have varied levels of success explaing sometimes contradicting psychophysical data for each stage (e.g. in the case of the existance of "half-axis" mechanisms); notably, in the case of color perception, as we move from the retina to higher order processing, the uncertainty of what mechanisms account for a large variety of results increases. In this dissertation, I use sophisticated psychophysical methods, combined with mathematical models of perception, to investigate long held assumptions about performance-based luminance. Specifically I test the plausibility of the assumption of coplanarity for colored lights, which its luminance is made equal to some given achromatic light using a task based on first order motion. I use the motion-equiluminant lights so obtained to investigate the mechanisms used by the human visual system to make judgments about motion-equiluminant colored textures.

In chapter one I describe a new method, based on minimum motion, to obtain motion-equiluminant lights, and I test the planarity of the motion equiluminant surface. This technique reveals properties of the motion-equiluminant plane that were previously undetected

using other methods. In particular, I find that the motion-equiluminant lights deviate from planarity, and do so following a pattern.

In chapter two I compare this new method to one used to derive photopic luminance, Heterochromatic Flicker Photometry (HFP). Using HFP I find that deviations of equiluminant lights from planarity are small, suggesting that the assumption of planarity is plausible. Yet, I argue that the experimental data from both techniques shows that the reason for this discrepancy is that minimum motion yields finer grained measurements compared to the measurements obtained using HFP.

Finally, in chapter three I model higher order visual mechanisms sensitive to color, rooted on the differential activation of each of the three cone classes. Using the minimum motion method presented in chapter one to acquire individualized motion-equiluminant color palettes, I apply the seed expansion method to create colored white noise textures. I find that performance in a detection task can be explained by two pairs of complementary full axis mechanisms: one with sensitivity that increases linearly across the gamut from red to green and one that increases linearly across the gamut from green to red, forming one complementary pair, and one with sensitivity that increases evenly with increased red-green saturation, and one that increases evenly with decreased red-green saturation, to form the other complementary pair. However, such explanation does not strongly preclude the existence of half-axis mechanisms.

Chapter 1

Refining the Minimium Motion Method for Equiluminance

1.1 Introduction

This chapter provides a new method for deriving performance-based equiluminant lights; specifically, we show that a motion-defined luminance setting provides an accurate (as suggested by confidence intervals smaller than those yielded by other performance-based methods) setting of lights motion-equiluminant to a given achromatic light. We explore its planarity, and in chapter 2 we compare this method with heterochromatic flicker photometry.

The minimum motion method for obtaining performance-based equliminant lights was introduced in the early 1980s as an improvement over heterochromatic flicker photometry and adjustment of borders (Anstis and Cavanagh, 1983), the fine tuning we describe is based on van Santen and Sperling (1984).

Even so, Heterochromatic Flicker Photometry (HFP) has been the standard for finding equiluminant lights that are performance-based in psychophysics research (Ives, 1912a,b,c,d,e, 1914, 1917), and is still used for that purpose.

One contention is that luminance, as defined by the CIE 85 years ago on 2 experiments using HFP, lacks a principled *raison d'être*, based as it is on the standard human observer, an averaged response from 17 subjects (10 from one experiment, 7 from another). We describe a modification of the minimum motion method that enables much more precise estimation of motion-defined luminance than has previously been available. Such modification and its concomitant increased sensitivity will prove crucial to reveal a previously hidden pattern of deviations from planarity, and will consequently challenge a long held notion about performance-based measures of luminance.

1.1.1 Testing the linearity of motion-luminance

As originally conceived by the International Commission on Illumination (Commission internationale de l'eclairage, CIE), the "photopic luminance" of a light with spectrum $Q(\lambda)$ is a statistic that is intended to reflect the effective intensity of the light for human vision. The photopic luminance of the light is computed by taking

$$\int \mathbf{V}(\lambda)Q(\lambda)d\lambda = \mathbf{V}^T Q \tag{1.1}$$

where $\mathbf{V}(\lambda)$ is the spectral luminous-efficiency function derived by the CIE in 1924 by combining data from seven studies, using different methods, but mainly Heterochromatic Flicker Photometry (HFP) and step-by-step brightness matching. Sharpe et al. (2005) proposed that the photopic luminance function be replaced by a new function that they call "$\mathbf{V}^\star(\lambda)$."

However, the basic idea of computing photopic luminance by taking the inner product of the light spectrum with some function of wavelength is retained in Sharpe et al. (2005).

Although there are many practical reasons why it might be useful to measure photopic luminance, there are also obvious reasons to doubt the theoretical justification for such a measure. In particular, human vision embodies a host of different processes whose effectiveness depends on light intensity. There is no good reason to suppose that the effective intensity of light (as a function of wavelength) should be the same for all of these processes. Indeed, as observed by Sharpe et al. (2005), "there are considerable differences between the luminous efficiency functions obtained by different measurement procedures and criteria, which include heterochromatic flicker photometry (HFP) or minimum flicker, a version of minimum flicker called heterochromatic modulation photometry (HMP), direct heterochromatic brightness matching, step-by-step brightness matching, minimally distinct border (MDB), minimum motion, color matching, absolute threshold, increment threshold, visual acuity, and critical flicker frequency." Moreover, for any one of these methods, different observers yield different luminous efficiency functions.

With that said, however, for any one of these methods m and any given participant, p, one might hope to be able to measure a luminous efficiency function $\mathbf{V}_{mp}(\lambda)$ analogous to $\mathbf{V}(\lambda)$. If so, then the luminance $\text{Lum}_{pm}(Q)$ for participant p of the light with spectrum Q measured using method m can be captured by

$$\int \mathbf{V}_{pm}(\lambda)Q(\lambda)d\lambda = \mathbf{V}_{pm}^T Q \tag{1.2}$$

for some function $\mathbf{V}_{pm}(\lambda)$ (which can be thought of as reflecting the luminous efficiency of different wavelengths for participant p in the context of method m), where the integral is over all wavelengths λ in the visible range. One goal of the current study is to determine whether luminance as measured using the minimum motion method (which we will call

"motion-luminance") (Anstis and Cavanagh, 1983) can be captured by Eq. 1.2 for some function $\mathbf{V}_{mp}$.

1.1.2 Eq. 1.2 implies that equiluminant lights should lie in a plane

Let f_1, f_2 and f_3 be the spectra of the three primaries used to generate the lights in a video monitor, and let B be a matrix whose three columns compose an orthonormal basis of the space Ω of functions spanned by f_1, f_2 and f_3 (noting that many of the functions in Ω will not be light spectra because they will take negative values). On the principal display device used for these experiments, the red, green and blue primaries took the forms shown in Fig. 1.1, and the three components of the basis B (which we think of a matrix with the three column vectors B_1, B_2 and B_3) that was derived for these spectra by the Matlab function orth and that was used are shown in Fig. 1.2.

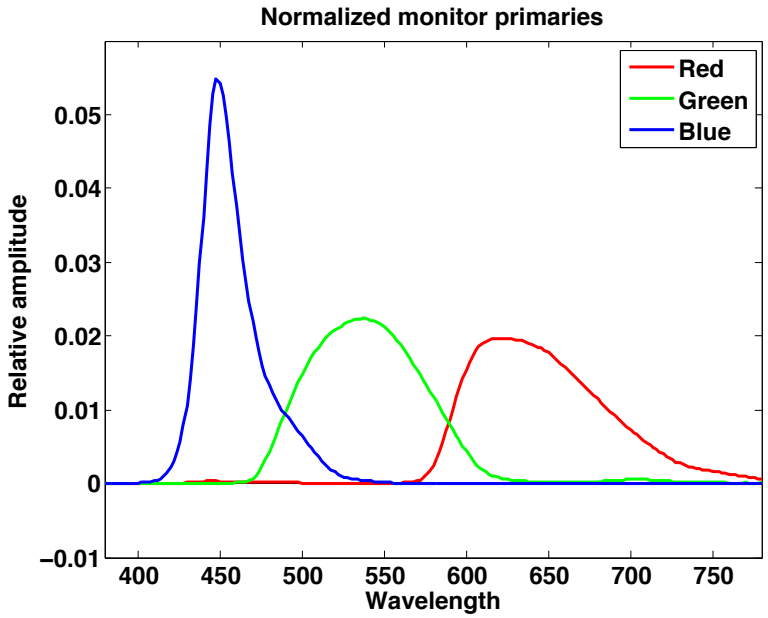

Figure 1.1: The red, green and blue primaries of the main monitor used.

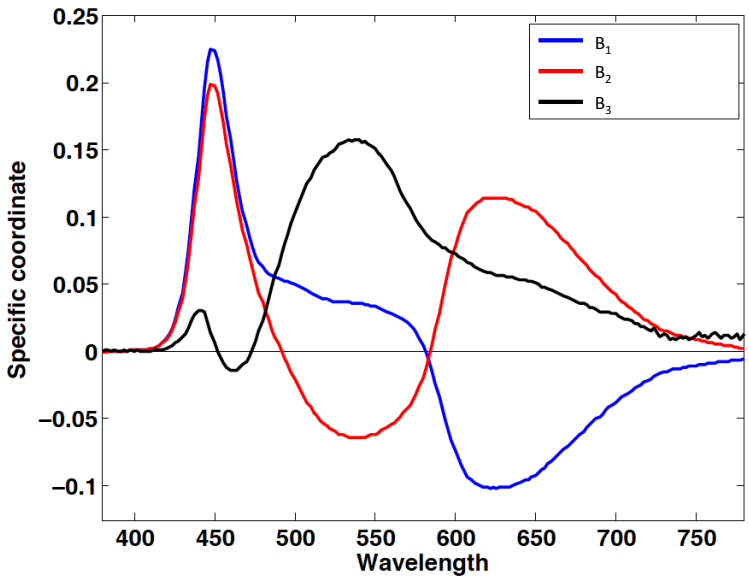

Figure 1.2: The orthonormal basis B of the space spanned by the primaries shown in Fig. 1.1 used to investigate the question of whether or not motion-equiluminance can be captured by Eq. 1.2.

Any function $Q \in \Omega$ corresponds to a unique vector $q \in \mathbb{R}^3$ such that

$$Bq = Q \qquad \text{or equivalently} \qquad q = B^T Q. \tag{1.3}$$

The three numbers in q are the linear-combination weights that must be given to the three column-vectors in B in order to derive the function Q. We call q the *code* of Q relative to the basis B, skipping the phrase "relative to basis B" whenever doing so will cause no confusion.

Suppose that the luminance of the light with spectrum Q measured for participant p using the minimum motion method is given by

$$\text{Lum}_p(Q) = \mathbf{V}_p^T Q \tag{1.4}$$

for some function $\mathbf{V}_p$. (From here on we will focus exclusively on luminance as measured using the minimum motion method and therefore drop the subscript "m" from our notation.)

There is no guarantee that the function $\mathbf{V}_p$ will be an element of Ω (i.e., that $\mathbf{V}_p$ can be expressed as a weighted sum of the three columns of B). The portion of the function $\mathbf{V}_p$ that *does* reside in Ω is called the *projection* of $\mathbf{V}_p$ into Ω. This is the function

$$V_p = Bv_p, \qquad \text{with code} \qquad v_p = B^T\mathbf{V}_p. \tag{1.5}$$

In general, $V_p \neq \mathbf{V}_p$; however, for any function Q that resides in Ω, it is only the projection of $\mathbf{V}_p$ into Ω that influences the inner product of $\mathbf{V}_p$ with Q. Thus,

$$\mathbf{V}_p^T Q = V_p^T Q. \tag{1.6}$$

where the 3-dimensional vectors v_p and q are the B-codes of the functions Q and V_p. In other words, if Q is an element of Ω, and if the motion luminance of Q is the inner product of $\mathbf{V}_p$ with Q, then the motion luminance of Q is equal to the inner product of V_p with Q. Thus, in particular, if Q is a light spectrum, and Eq. 1.4 holds, then

$$\mathbf{V}_p^T Q = V_p^T Q = (Bv_p)^T Bq = v_p^T B^T Bq = v_p^T q, \tag{1.7}$$

showing that for participant p the motion-luminance of the light with spectrum Q is the inner product of the B-codes $v_p^T q$.

Note, moreover, that if Eq. 1.4 holds, then the set of all B-codes q corresponding to the spectra $Q \in \Omega$ of lights that are motion-equiluminant for participant p to a given light Q_0 with code q_0 satisfies

$$v_p^T q = v_p^T q_0 \tag{1.8}$$

implying that $v_p^T(q - q_0) = 0$, i.e., the vector difference between q and q_0 is perpendicular to v_p. Thus, if we write q as $q_0 + (q - q_0)$, we see that all of the codes q corresponding to lights

equiluminant to the light with spectrum Q_0 lie in the plane that contains the code q_0 and is perpendicular to v_p.

One purpose of the current experiment is to test this prediction.

1.1.3 Motion-equiluminance

In the experiment described below, we obtain 20 colors, varying in hue, each of which is equated in motion-luminance to a fixed gray. Each of the colored lights is derived using a fine tuning of the minimum motion method. The method was introduced first by Anstis and Cavanagh (1983), the fine tuning is based on Lu and Sperling (1995) who found that luminance grating contrasts of 4-8% provided maximum sensitivity for producing equiluminant red-green gratings; we use 6% contrast modulation trhoughout.

The stimulus we use is illustrated in Fig. 1.3. The stimulus is periodic in time with four frames, each of which comprises an annular section of a radial square-wave grating that runs through four cycles per circumference. Frame 1 is shown on the left of Fig. 1.3. In Frame 2, the square wave has been shifted one quarter cycle; if we think of this shift as clockwise (which corresponds to downward on the right side of Fig. 1), then the sections that were dark in Frame 1 have become green, and the sections that were light in Frame 1 have become gray. Frame 3 is identical to Frame 1, except that the dark and light bars switch roles, and Frame 4 is identical to Frame 2, except that the green and gray bars switch roles.

The method assumes that the motion of this stimulus is registered by only a single motion-sensing system in human vision. Although there is overwhelming evidence to suggest that several different systems operate in human vision to detect motion (Lu and Sperling, 1995), the spatiotemporal properties of this stimulus (details described below) are chosen to acti-

vate exclusively the "first-order" system, i.e., the system that is sensitive to spatiotemporal modulations of light intensity.

The logic of the minimum motion method is suggested by the light and dark arrows on the right side of Fig. 1. We assume that each of the four lights that occur in this stimulus (the dark-gray, the light-gray, the middle-gray and the green), produces a level of activation that we will call "motion-luminance" in the unique system that registers the motion of this stimulus, and spatiotemporal variations in motion-luminance are analyzed for motion. In the following two paragraphs, imagine that the four frames of the stimulus diagrammed in Fig. 1.3 are repeated a number of times at a high temporal frequency.

Suppose that the motion-luminance of the green light is higher than that of the middle-gray light. In this case, all standard models of first-order motion perception (Adelson and Bergen, 1985; van Santen and Sperling, 1984, 1985; Watson and Ahumada, 1985) predict that the motion evoked by the stimulus of Fig. 1.3 should follow the upward/counterclockwise path suggested by the white arrow in the right side of Fig. 1.3. Conversely, if the motion-luminance of the green light is lower than that of the middle-gray light, the motion evoked should follow the downward/clockwise path suggested by the black arrow.

By contrast, suppose the motion-luminance of the green light is equal to that of the middle-gray light. In this case, stimulus frames 2 and 4 become spatially uniform in motion-luminance, in which case the stimulus devolves into pure flicker. The homogeneous motion-luminance annuli that now occur in frames 2 and 4 serve merely to mark time between the patterns of motion-luminance occurring in frame 1 and 3; however, the pattern of motion-luminance in frame 3 is merely the contrast-reversal of the pattern in frame 1, implying that this stimulus must evoke completely ambiguous motion.

These observations suggest a natural strategy for finding lights that are equiluminant to a given standard light. Suppose that lights Q_i increase by small steps in motion-luminance

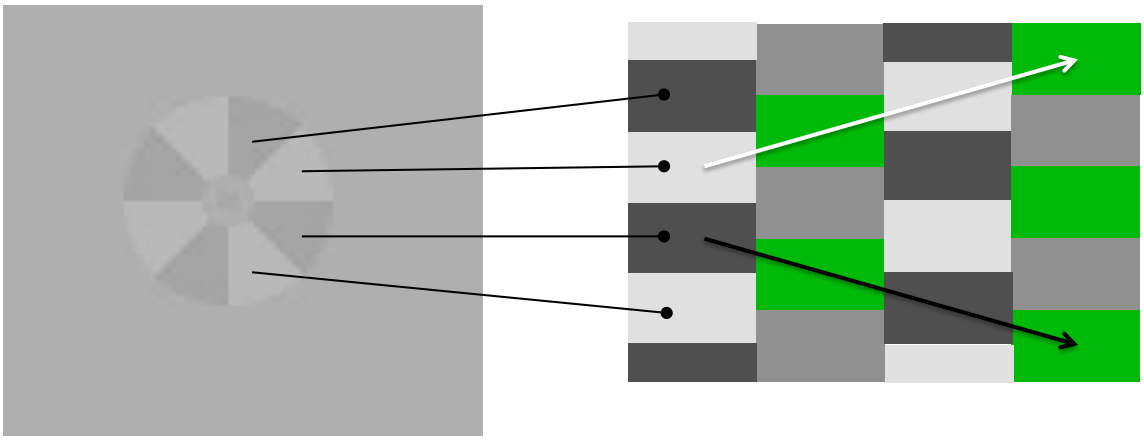

Figure 1.3: On the left, a static depiction of the moving stimuli. On the right, a vertical representation of the radial stimulus. Each colum represents a frame (four vertical frames are depicted). Each frame is replaced by the adjacent, occupying the same location, so that a slight displacement of the gating is perceived as the frames are rapidly presented.

as i ranges from 1 to N. Then engage the participant in a task requiring him/her to judge the direction of motion evoked by the stimulus diagramed in Fig. 1, using a 1-up-1-down staircase to control the index i of the light presented on each trial. That is, on a given trial, if the stimulus presented used light Q_i in the role of the green in Fig. 1.3, and participant's response suggests that the motion-lumimince of the Q_i is higher (lower) than the standard light (the medium-gray light in Fig. 1.3), then increment (decrement) i on the next trial. One can fit a psychometric function to the resulting data (e.g., a cumulative normal distribution function) and take the 50% threshold as an estimate of the equiluminant point. This is precisely the strategy we will use below.

1.2 Methods

All methods were approved by the UC Irvine Institutional Review Board, and all participants provided signed consent.

1.2.1 Participants

All 9 participants (four female) reported normal or corrected-to-normal vision; color vision was assessed using Ishihara plates and was found to be normal.

1.2.2 Apparatus

Stimuli were presented using the Psychophysics Toolbox (Brainard, 1997) running under Matlab on an iMac monitor (iMac 21' 2009 late model). The subjects sat approximately 90 cm from the screen without a chin rest. The room was completely darkened and the only source of illumination was the monitor.

1.2.3 Procedure

The display device we used afforded 8-bit resolution in red, green and blue primaries. Thus, the achievable lights corresponded to vectors $(r, g, b)^T$ in which each of r, g and b was an integer between 0 and 255. We will relax this constraint, however, and call vector $v \in \mathbb{R}^3$ a "pixel value" if all of r, g and b are greater than or equal to 0 and less than or equal to 255. Any pixel value whose coordinate values are all integers will be called "achievable" to indicate that we can actually produce the light to which it corresponds. We use the transpose sign in writing "$(r, g, b)^T$" because for most of the computations we will want to perform it will be convenient to think of pixel values as column vectors.

Our goal is to derive for each participant a set of 20 lights varying widely in hue, each of which is motion-equiluminant to the fixed standard light G with pixel value $v_G = (175, 175, 175)^T$. This light is achromatic and has photometric luminance $52 cd/m^2$, which is roughly half the

luminance of the light with pixel value $(255, 255, 255)^T$, the brightest achievable achromatic light.

On the monitor we were using, previous work suggested that the luminances of the red, green and blue primaries were roughly proportional to $1.5 : 4 : 0$. Accordingly, for a vector $v = (1.5, 4, 0)^T$, we took the direction $\hat{l} = \frac{v}{\|v\|} = (0.3511, 0.9363, 0)^T$ as a rough approximation of the luminance axis in the space of lights spanned by our red, green and blue primaries. Note that for $\hat{l}_1$ and $\hat{l}_2$ the first and second coordinate values of $\hat{l}$, the columns of the matrix

$$R = \begin{bmatrix} 0 & -\hat{l}_2 \\ 0 & \hat{l}_1 \\ 1 & 0 \end{bmatrix} \tag{1.9}$$

compose an orthonormal basis of the space of RGB triples orthogonal to $\hat{l}$. We proceeded to derive the 10 vectors

$$\phi_k = R \begin{bmatrix} \cos\rho_k \\ \sin\rho_k \end{bmatrix} \qquad \text{for} \qquad \rho_k = \frac{\pi}{20} + \frac{k\pi}{10}, \qquad k = 0, 1, \cdots, 9. \tag{1.10}$$

The vectors ϕ_k have norm 1, and they are spaced at equal angles in the plane orthogonal to $\hat{l}$.

For a given $k = 0, 1, \ldots, 9$, the plane of lights with pixel values expressible as linear combinations of $\hat{l}$ and ϕ_k includes exactly two maximally saturated lights that are equiluminant with the standard light G. We will use the minimum motion technique to estimate each of these two lights. The first step is to define what we call "the perimeter of the plane spanned by ϕ_k and $\hat{l}$." This is the set of pixel values corresponding to the most extreme lights that can be generated by deviating away from v_G (the pixel value of G) in the plane spanned by ϕ_k and $\hat{l}$. For any angle ρ, we can derive the pixel value v_ρ in this set by heading out away

from v_G in the plane with x-axis ϕ_i and y-axis $\hat{l}$ in the direction with polar angle ρ. As we head out away from v_G in this particular direction, eventually we run out of room when one of the three coordinate values of the vector we are extending becomes either 0 or 255. Formally,

$$v_\rho = v_G + \alpha_\rho \left(\cos(\rho)\phi_k + \sin(\rho)\hat{l} \right), \tag{1.11}$$

where α_ρ is the smallest scalar such that either (1) the minimum coordinate value of v_ρ is zero or else (2) the maximum coordinate value of v_ρ is 255.

Of course, the perimeter P_k of the plane spanned by ϕ_k and $\hat{l}$ contains many pixel values that are not achievable because one or more of their coordinate values fail to be integers. We use the following set of 400 pixel-values to approximate P_k:

$$w_{\rho_h} = \text{round}(v_{\rho_h}) \qquad \text{where} \qquad \rho_h = \frac{2\pi h}{400} \quad \text{for} \quad h = 0, 1, \cdots, 399. \tag{1.12}$$

The perimeters of the planes spanned by ϕ_k and $\hat{l}$, for $k = 0, 1, \cdots, 9$ are shown in Fig. 1.4. For ease of reference, we also introduce table 1.1 with the nomenclature that will be used throughout this and the following chapter.

Each of the perimeters P_k, $k = 0, 1, \cdots, 9$ intersects the surface of equiluminant lights at two points. We use the minimum motion method to estimate each separately. To estimate one of them, the participant began by adjusting the color of a square subtending ≈ 0.5 deg. visual angle, placed in the center of the screen, to be roughly equiluminant with a homogeneous background of standard light G. By repeatedly pressing the left-arrow (right-arrow) key, the participant could move around P_k in one (the other) direction.

When the subject judged that the colored spot was nearly equiluminant to the background, he/she pressed the "o" key. The point of this initial adjustment procedure was to focus

Name	Angle	R	G	B	B-code 1	B-code 2	B-code 3	x	y	Y
R	-176	255	135	161	-0.00256	0.00825	0.00689	0.431	0.305	0.144
O	-163	255	136	130	-0.00354	0.00729	0.00680	0.465	0.341	0.143
O	-152	255	145	90	-0.00439	0.00639	0.00699	0.499	0.386	0.148
O	-146	255	148	10	-0.00508	0.00574	0.00704	0.533	0.426	0.149
Y	-111	206	179	0	-0.00245	0.00274	0.00713	0.453	0.488	0.159
YG	-77	152	193	0	-0.00046	0.00046	0.00668	0.385	0.542	0.154
G	-59	90	202	0	0.00091	-0.00112	0.00657	0.327	0.588	0.156
G	-54	1	209	0	0.00154	-0.00184	0.00673	0.300	0.610	0.161
G	-49	0	207	86	0.00212	-0.00126	0.00660	0.283	0.547	0.158
A	-33	0	203	148	0.00365	0.00011	0.00659	0.253	0.434	0.157
T	0	0	198	202	0.00592	0.00209	0.00664	0.225	0.334	0.157
B	38	13	186	255	0.00892	0.00495	0.00613	0.199	0.240	0.143
B	44	91	176	255	0.00822	0.00568	0.00592	0.212	0.230	0.136
L	55	129	168	255	0.00749	0.00657	0.00588	0.228	0.221	0.133
L	66	157	162	255	0.00674	0.00742	0.00605	0.245	0.219	0.135
L	80	180	153	255	0.00595	0.00833	0.00600	0.261	0.211	0.131
V	98	206	142	255	0.00489	0.00954	0.00616	0.283	0.206	0.132
V	120	239	120	255	0.00323	0.01145	0.00629	0.315	0.196	0.130
V	145	255	121	223	0.00049	0.01075	0.00666	0.361	0.229	0.138
R	168	255	130	189	-0.00126	0.00915	0.00680	0.397	0.271	0.142

Table 1.1: Columns are: Common color name (R=red, G=green, B=blue, Y=yellow, O=orange, YG=yellow-green, A=aqua, T=turquoise, L=lavender, V=violet), Hue Angle, RGB pixel values, B-codes (1,2,3), and CIE 1931 xyY values.

all of the trials in the motion-direction task (described below) in the neighborhood of the equiluminant light to be estimated.

Following this initial, rough adjustment procedure, the subject performed 60 trials in the minimum motion task (for a few subjects 120 trials were needed in some conditions). On a given trial, the subject fixated a small cue spot in center of the screen and pressed the space bar to initiate a trial. Eight frames of the annular motion stimulus were then displayed at the rate of 60 frames per second, producing a stimulus temporal frequency of 15 Hz. The outer and inner diameters of the motion annulus were 1.24 and 0.28 deg. respectively. The stimulus regions colored medium gray in Fig. 1.3 were populated with standard light G. The stimulus regions colored light gray and dark gray in Fig. 1.3 were achromatic with Weber contrasts

0.06 and -0.06 respectively relative to G. The stimulus regions colored green in Fig. 1.3 were populated with a test light T that was varied from trial to trial. The temporal frequency of 15 Hz was high enough to exclude any contribution from the third-order motion system, and the small deviations in Weber contrast of the light and dark grays in the stimulus from G excluded potential contributions from the second-order motion system (Lu and Sperling, 1995). Moreover, it has been shown that the sensitivity of of minimum motion method to deviations in motion-luminance of the test light from the standard is maximized by using low Weber contrasts (around ± 0.06) in the light and dark gray lights that occur in alternate frames (Lu and Sperling, 2001).

After the stimulus was presented, the subject pressed the right-arrow key if the motion appeared clockwise or the left-arrow key if counterclockwise. A 1-up-1-down staircase procedure was used to concentrate observations in the neighborhood of the light that was motion-equiluminant to G. Specifically, if the current trial used a test light T (to populate the regions in Fig. 1.3 colored green) with pixel value (r, g, b), then if the participant's response indicated that T was higher (lower) in motion-luminance than G, the next test-light pixel value was the nearest one to (r, g, b) in the current perimeter lower (higher) in $\hat{l}$.

This same procedure was repeated twice to derive 60 (in a few cases 120) trials of data on each side of each of the 10 perimeters P_k. Thus, the data derived from a given participant comprised 20 sets of psychometric data (one set for each of the 20 maximally saturated equiluminant color estimates).

1.2.4 Modeling

For any pixel value v, let $\mathrm{Lum}_p(v)$ be the motion-luminance of the light generated by pixel value v for participant p. Each of the 20 staircases yielded data that comprised

1. the list of pixel values v_j that was used to generate the test light on one or more trials, where $\hat{l}^T v_j$ is a increasing function of the index j;

2. the number K_j of trials on which v_j was presented and the participant's response indicated that $\text{Lum}_p(v_j) > \text{Lum}_p(v_G)$, and

3. the number N_j of trials on which v_j was presented and the participant's response indicated that $\text{Lum}_p(v_j) < \text{Lum}_p(v_G)$.

1.2.4.1 Fitting the data from a single staircase.

Writing Φ for the standard normal cumulative distribution function, we assume that across the range of pixel values v_j visited by the staircase,

$$\mathbb{P}\left[\text{Lum}_p(v_j) \text{ judged higher than } \text{Lum}_p(v_G)\right] = \Phi\left(\frac{j-\mu}{\sigma}\right) \tag{1.13}$$

where the mean μ is the real number (sampled by the integer values j used to index the actual pixel values tested) for which $\text{Lum}_p(v_\mu) = \text{Lum}_p(v_G)$.

The likelihood function for the parameters μ and σ is thus given by

$$\Lambda(\mu,\sigma) = \prod_{\text{indices } j} \Phi\left(\frac{j-\mu}{\sigma}\right)^{K_j} \left(1-\Phi\left(\frac{j-\mu}{\sigma}\right)\right)^{N_j}. \tag{1.14}$$

We assume a uniform prior on μ and σ and proceed to use Markov chain Monte Carlo simulation to derive a sample of pairs (μ_i, σ_i), $i = 1, 2, \cdots, 8000$, from the joint posterior distribution characterizing μ and σ.

Each of the 8000 resulting values of μ_i is then converted back into a B-code m_i via linear interpolation as follows. Let j be the index such that $j < \mu_i < j+1$, and let q_j and q_{j+1} be

the B-codes of the lights with pixel values v_j and v_{j+1}. We take

$$m_i = \rho q_{j+1} + (1-\rho)q_j \quad \text{for} \quad \rho = \mu_i - j. \tag{1.15}$$

Thus we derive 8000 estimates m_i of the B-code for which the light Bm_i is motion-equiluminant to G.

1.2.4.2 Estimating the motion-equiluminant plane and motion-luminance axis.

For $i = 1, 2, \cdots, 8000$, let m_{ik} be the B-code derived from the i^{th} iteration of the Markov chain Monte Carlo process for the k^{th} set of psychometric data. For each $i = 1, 2, \cdots, 8000$, we will extract from the 20 B-codes, $m_{i1}, m_{i2}, \cdots, m_{i20}$ an estimate of the motion-equiluminant plane for participant p (as well as of the motion-luminance axis, l_p, for participant p). To do this, we proceed as follows:

1. Let M_i be the mean of the 20 B-codes $m_{i1}, m_{i2}, \cdots, m_{i20}$, and set

$$\delta_{ik} = m_{ik} - M_i \quad \text{for} \quad k = 1, 2, \cdots, 20. \tag{1.16}$$

2. Then for Δ_i the 3×20 matrix whose k^{th} column is δ_{ik}, use singular value decomposition to derive

$$C_i \times E_i \times L_i^T = \Delta_i \tag{1.17}$$

for C_i a 3×3 orthonormal matrix (of principal components), E_i a nonnegative diagonal matrix (of eigenvalues) with entries decreasing downward, and L_i a 20×3 column orthonormal matrix (of component loadings).

3. The least-squares estimate of the motion-equiluminant plane is then (using Matab-like notation) the set of all B-codes

$$q_v = M_i + C_i(:, 1:2) \times v \tag{1.18}$$

where $C_i(:, 1:2)$ is the 3×2 matrix comprising the first two columns of C_i, and $v_q \in \mathbb{R}^2$. In particular, the projections of the B-codes m_{ik}, $k = 1, 2, \cdots, 20$ into the motion-equiluminant plane are given by

$$\hat{m}_{ik} = M_i + C_i(:, 1:2) \times E_i(1:2, 1:2) \times L_i(:, 1:2)^T, \tag{1.19}$$

where $C_i(:, 1:2)$ is the 3×2 matrix comprising the first two columns of C_i, $E_i(1:2, 1:2)$ is the 2×2 matrix comprising the first two rows and columns of E_i and $L_i(:, 1:2)$ is the 20×2 matrix comprising the first two columns of L_i.

4. The least-squares estimate of the motion-luminance axis derived from the i^{th} iterations of all 20 Markov chain Monte Carlo simulations is the unit vector $l_i = C_i(:, 3)$, the third column of the matrix C_i.

1.2.4.3 In summary

We use Markov chain Monte Carlo simulation to derive a sample of size 8000 drawn from the posterior joint density characterizing

1. each of the B-codes m_k, $k = 1, 2, \cdots, 20$;
2. the motion-equiluminant plane (i.e., the plane such P that the sum of squared Euclidean distances of the B-codes m_k, $k = 1, 2, \cdots, 20$, from P is minimal);

3. the motion-luminance axis l (i.e., the unit B-code orthogonal to the motion-equiluminant plane).

From these various samples, we derive the 95% credible intervals used in all the different figures.

1.3 Results

The raw pixel values of equliluminant settings obtained for 9 participants are shown in Fig. 1.4. Each black circle shows the estimated pixel value along a given hemi-perimeter for a given participant. Pixel values, however, are related to spectral variations by nonlinear transformations that make the results shown in Fig. 1.4 difficult to interpret.

In Fig. 1.5 are plotted the B-codes of the 20 lights estimated, for one participant p (additional results for whom are plotted in panel 1 of Fig. 1.7), to be motion-equiluminant to the light G with pixel value $v_G = (175, 175, 175)^T$ whose B-code we denote q_G. Fig. 1.6 plots the 1931 CIE x and y coordinate values of the 20 lights. In order to facilitate comparison of this figure with Fig. 1.7, this figure also labels several of the points in this plot with the hue angles used to label them on the horizontal axis in Fig. 1.7.

Although the colored circles seem close to the best-fitting plane, in most cases, they deviate significantly from this plane. This is shown by Fig. 1.7 which plots the deviations of the 20 B-codes corresponding to the motion-equiluminant lights from the best-fitting planes estimated individually for all 9 participants.

The ordinate of this plot needs some explanation. The plane that minimizes the sum of squared distances to the 20 B-codes q corresponding to lights Q estimated to be equiluminant

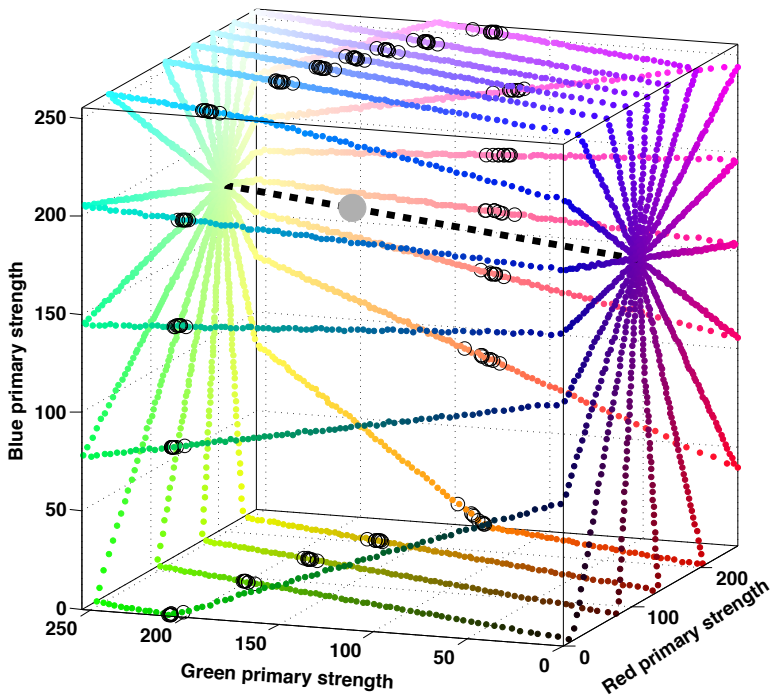

Figure 1.4: *The 10 perimeters of the planes spanned by ϕ_k and $\hat{l}$ for $k = 0, 1, \cdots, 9$.* The large gray disk indicates the pixel value $v_G = (175, 175, 175)^T$ used to generate the standard light G. The dashed black line passing through v_G represents the locus of pixel-values $v_G + \alpha\hat{l}$ for $\alpha \in \mathbb{R}$. A given perimeter corresponds to the intersection of the plane that can be generated by taking $v_G + \alpha\hat{l} + \beta\phi_k$ for $\alpha, \beta \in \mathbb{R}$, $k = 0, 1, 2, \cdots, 9$, with the surface of the cube of all possible pixel values.

for participant p is defined by

$$l_p^T q = l_p^T \mu_p \tag{1.20}$$

for the normalized B-code l_p (i.e., $\|l_p\| = 1$) orthogonal to the best-fitting plane and μ_p the mean of the 20 B-codes q. For any B-code q, we can think of $l_p^T q$ as the activation produced in a mechanism with sensitivity function $L_p = Bl_p$ by the light $Q = Bq$. To see this, note that

$$l_p^T q = l_p^T B^T B q = (Bl_p)^T Bq = L_p^T Q. \tag{1.21}$$

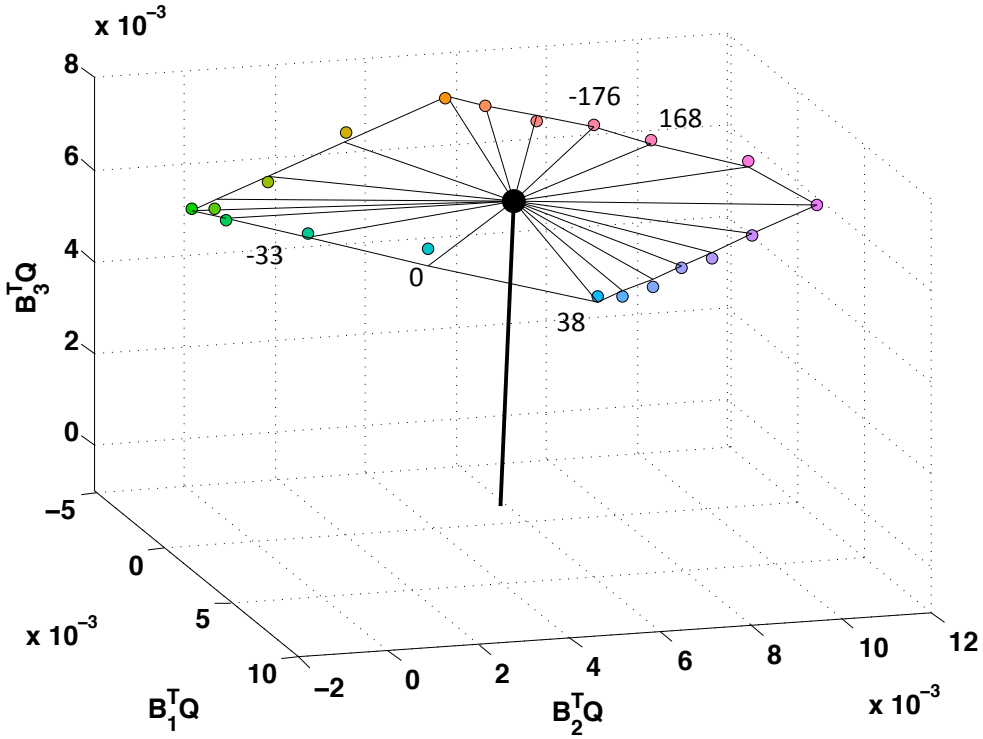

Figure 1.5: *The 20 B-codes corresponding to the motion-equiluminant lights estimated for one participant p.* The thick black bar is oriented in the direction of the B-code of the estimated motion-luminance axis l_p for this participant with length scaled to be equal to $l_p^T q_G$, the B-code of the standard light G with pixel value $v_G = (175, 175, 175)$. The thin black lines all lie in the estimated motion-equiluminant plane which is the plane of B-codes q satisfying $l_p^T q = l_p^T q_G$. The colored circles show the B-codes of the individually estimated, motion-equiluminant lights. The numbers show, for some of the B-codes, color angles (from table 1.1) associated with deviations from the best-fitting plane plotted in Fig. 1.7.

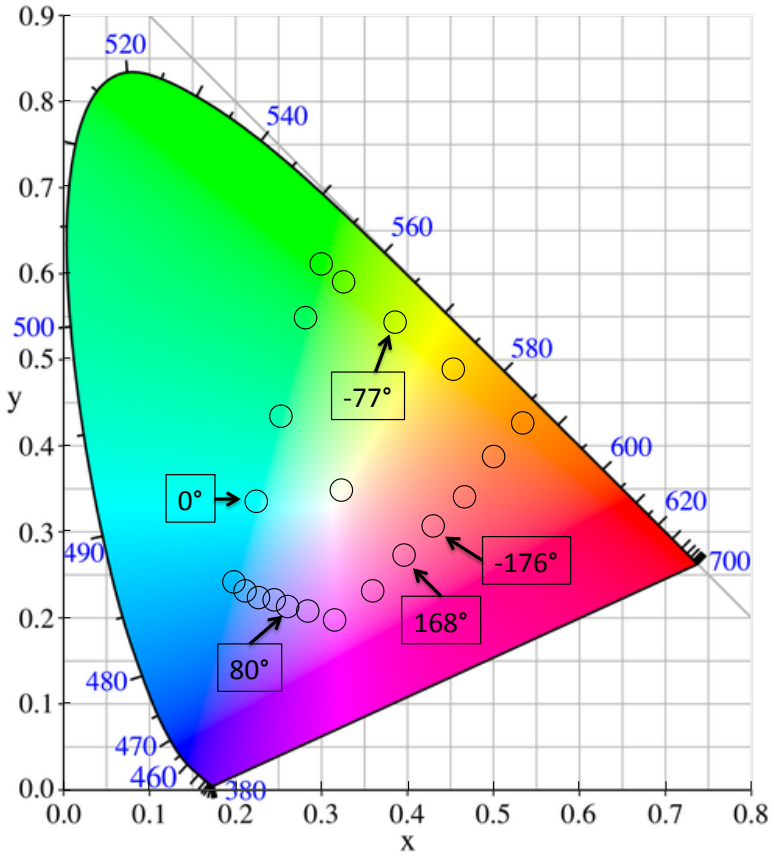

Figure 1.6: *The CIE 1931 x and y values corresponding to the hues of the motion-equiluminant lights estimated for the same participant whose results are plotted in 1.5.* The numbers associated with some of the hues are the hue angles on the horizontal axis in Fig. 1.7. The circle in the center corresponds to the standard gray to which lights were made equiluminant.

If in fact motion-luminance satisfied Eq. 1.1, then for any light Q with B-code q, $\mathrm{Lum}_{pm}(Q)$ would be exactly equal to $l_p^T q$. However, the 20 B-codes q estimated to be motion-equiluminant for participant p do not lie exactly in a plane. To quantify the degree to which these B-codes q deviate from the best fitting plane plane, we take

$$\frac{l_p^T q}{l_p^T \mu_p} \tag{1.22}$$

as the ordinate in Fig. 1.7.

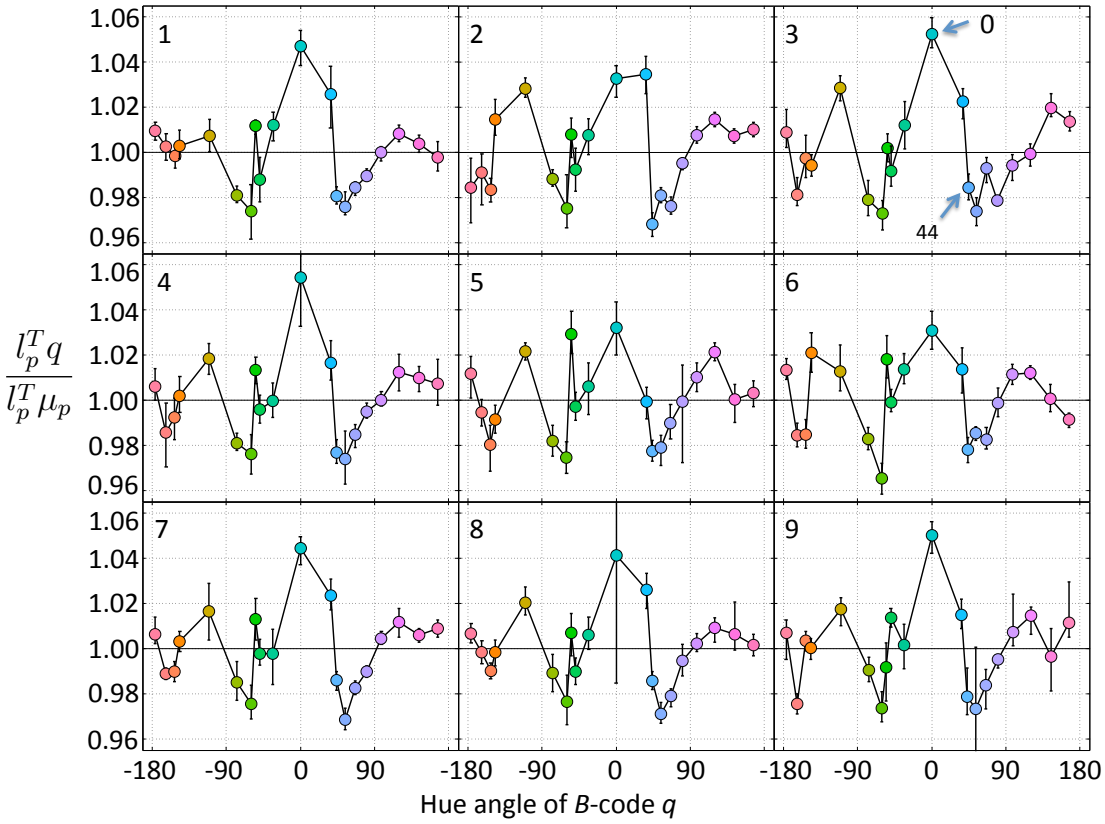

Figure 1.7: *The mean deviations of the 20 B-codes corresponding to the motion-equiluminant lights from the best-fitting plane estimated for 9 participants.* The error bars are 95% Bayesian credible intervals. For μ_p the mean of all 20 B-codes estimated for participant p, and l_p the normal to the best-fitting planar approximation to the 20 B-codes, the ordinate in the plot for participant p is $\frac{l_p^T q}{l_p^T \mu_p}$ for each of the 20 B-codes q estimated for participant p.

For Q the light with B-code q and M_p the light with B-code μ_p, Eq. 1.22 reflects the proportional deviation of the optimal linear approximation of $\text{Lum}_p(Q)$ to the optimal linear approximation of $\text{Lum}_p(M_p)$. Thus, if motion-luminance satisfied Eq. 1.1, $\frac{l_p^T q}{l_p^T \mu_p}$ would be equal to 1 for all 20 B-codes. Fig. 1.7 shows that this is not the case. The error bars are 95% Bayesian credible intervals. As is clear, most of these credible intervals do not contain 1. Moreover, the patterns of deviations are similar across participants.

Fig. 1.8 shows the average of the 9 curves plotted separately in Fig. 1.7. In this figure, the error bars are 95% confidence intervals derived using a t-distribution with 8 degrees

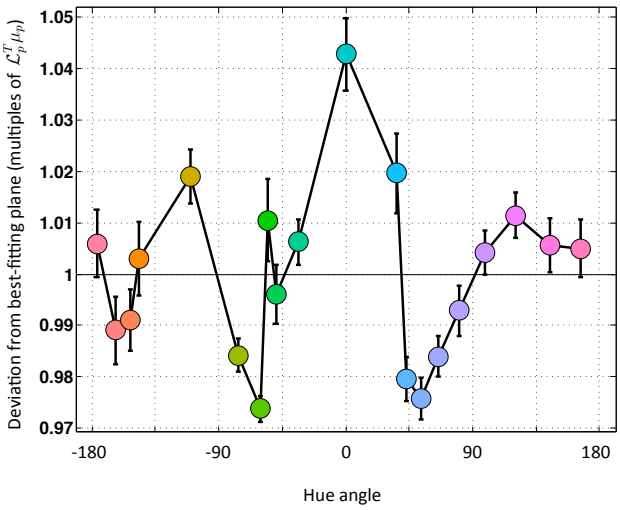

Figure 1.8: *The deviations of the 20 B-codes corresponding to the motion-equiluminant lights from the best-fitting plane averaged across all 9 participants.* The error bars are 95% confidence intervals derived using a t-distribution with 8 degrees of freedom.

of freedom. As is clear, for q_0 the B-code of the blue-green light assigned to hue angle 0 deviates strongly and positively from $l_p^T \mu_p$. This tells us that in order to be made motion-equiluminant to G, the spectrum of a light at hue angle 0 must be elevated in its amplitude by more than 4% relative to the planar prediction in which $l_p^T q_0 = l_p^T \mu_p$. By contrast, the B-codes of the slightly pinker blues with hue angles 44 and 55 deviate strongly and nagatively from $l_p^T \mu_p$. This tells us that in order to be made motion-equiluminant to G, the spectrum of a light with, for example, B-code q_{44}, with hue angle 44 must be *decreased* in its amplitude by slightly more than 2% relative to the planar prediction in which $l_p^T q_{44} = l_p^T \mu_p$.

1.3.1 The motion-luminance axes of our participants

The plane that minimizes the sum of squared distances to the 20 spectra Q estimated for participant p to be motion-equiluminant to the standard light is given by Eq. 1.20. However,

the function

$$L_p = Bl_p \tag{1.23}$$

is a linear combination of the red, green and blue primaries that happen to be resident in the monitor used in this experiment. Because these three functions need not reside in the space spanned by the human cone sensitivity functions, there may well be a component of L_p to which human vision is blind. Our aim in this section is to focus on the portion of L_p to which human vision IS sensitive, the projection of L_p into the space Ω spanned by the short, medium and long wavelength human cone fundamentals f_S, f_M and f_L. For F an orthonormal basis of Ω, the projection of L_p into Ω is

$$\mathcal{L}_p = FF^T L_p. \tag{1.24}$$

The projection $\mathcal{L}_p$ is invariant with respect to the particular orthonormal basis F that is used to generate it. We call this function the motion-luminance axis for participant p. In order to characterize the functions $\mathcal{L}_p$ estimated for our 9 participants p, we will use a basis derived from the Stockman-Sharpe 2 deg. cone fundamentals. Let $\tilde{f}_S$, $\tilde{f}_M$ and $\tilde{f}_L$ be the short-, medium- and long-wavelength Stockman-Sharpe cone fundamentals, and set

$$f_S = \frac{\tilde{f}_S}{\|\tilde{f}_S\|}, \qquad f_M = \frac{\tilde{f}_M}{\|\tilde{f}_M\|}, \qquad \text{and} \qquad f_L = \frac{\tilde{f}_L}{\|\tilde{f}_L\|}. \tag{1.25}$$

Then sequentially set

$$F_{M+L} = \frac{f_M + f_L}{\|f_M + f_L\|}, \tag{1.26}$$

$$F_S = \frac{h_S}{\|h_S\|}, \quad \text{for} \quad h_S = f_S - (f_S^T F_{M+L})F_{M+L}, \tag{1.27}$$

and

$$F_{L-M} = \frac{h_{L-M}}{\|h_{L-M}\|}, \quad \text{for} \quad h_{L-M} = f_L - f_M - (F_{M+L}^T(f_L - f_M))F_{M+L} - (F_S^T(f_L - f_M))F_S. \tag{1.28}$$

It is easy to check that F_{M+L}, F_S and F_{L-M} compose an orthonormal basis of the space spanned by the Stockman-Sharpe cone fundamentals. Let F be the matrix whose three columns are the functions F_{M+L}, F_S,and F_{L-M}. For any function Q in the space spanned by the human cone fundamentals, we call the column vector $F^T Q$ the F-code of Q.

In Fig. 1.9 are plotted the F-codes of the motion-luminance axes $\mathcal{L}_p$ of our 9 participants p. Panel A. shows all three dimensions of each F-code with the gray disk giving the mean F-code. Panel B shows the projections of the F-codes into the F_S, F_{L-S} plane. The numbers associated with different F-codes correspond to the numbers assigned to different panels in Fig. 1.7. There are several things to note:

1. By far the dominant F-code component for all participants p is $F_{M+L}^T \mathcal{L}_p$.
2. The average value of $F_S^T \mathcal{L}_p$ is near 0; however, $F_S^T \mathcal{L}_p$ differs significantly from 0 for nearly all individual participants p.
3. For all participants, $F_{L-M}^T \mathcal{L}_p$ is slightly negative.

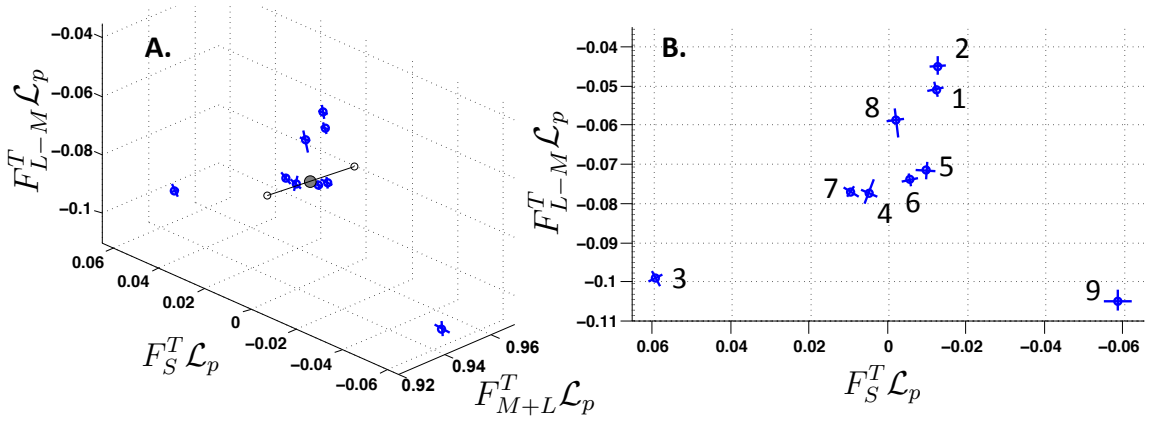

Figure 1.9: *The F-codes of the motion-luminance axes $\mathcal{L}_p$ of our 9 participants p.* A. All three dimensions are shown. The blue points give the motion-luminance axis F-codes for different participants. error bars are 95% Bayesian credible intervals aligned with the principal components of posterior density. The gray disk gives the average motion-luminance axis F-code, and the black line running through the gray disk indicates the mean direction out from the origin of the F-codes. Note that this mean direction is nearly parallel to the F_{M+L} axis. B. The projection of the motion-luminance axis F-codes into the F_S, F_{L-S} plane. The numbers associated with different F-codes correspond to the numbers assigned to different panels in Fig. 1.7.

1.4 Discussion for Chapter 1

Precise methods and instruments to measure spectral power distribution from an emitting source are readily available (eg. using a spectroradiometer). However, that measure does not provide much insight to our understanding of how the human visual system process that information. The CIE, nearly a hundred years ago, made an attempt to incorporate the effective intensity of light on human vision, through the computation of the photopic luminance, based on the spectral luminous efficiency function, itself derived mostly on heterochromatic flicker photometry and distinct border methods. We take the basic idea of computing photopic luminance by taking the inner product of the light spectrum with some function of wavelength, for a specific method, namely, a fine-tuned version of the minimum motion procedure (Anstis and Cavanagh, 1983; van Santen and Sperling, 1984; Lu and Sperling, 1995). Our version of the minimum motion method is modified in two important ways: although we kept the set up of interleaving alternating two square waves, one with dark gray/light

gray and other with target gray/target color, the frequency of the apparent circular motion was high enough to prevent contribution from the thrid-order motion system. Secondly, very small deviations on the Weber contrast of the section of our stimulus that "sandwiched" the color being adjusted for, both maximized the sensitivity to deviations to motion-luminance (Lu and Sperling, 2001), and excluded potential contributions from the second-order motion system (Lu and Sperling, 1995).

The underlying assumption of planarity was not held. Our results show that when using this minimum motion method, modified in the ways described, we obtain unprecedently precise measurements. What might have looked like a plane of equiluminant colors using other methods, under our method it is revealed to be an undulating ring of colors that deviate from planarity in a stable, predictable pattern. We are showing the results of the maximum saturation achievable in the monitor we used. We conjecture that exploring different saturation levels of motion-equiluminant lights will be revealed to be an undulating equiluminant surface.

Consequently, our results suggest that motion-luminance cannot be modeled as the inner product some spectra with some function of wavelength, such as the suggested $\mathbf{V}(\lambda)$. Specifically, the results sumarized on Fig. 1.8 lights that are perceptually equiluminant to a standad can still deviate from each other by more than 6% from each other, and as much as 4% from the best fitting plane. These results alone precludes the possibility of using the inner product model for motion-luminance.

Our results are in agreement with previous research on performance-based luminance, and show that motion-luminance is dominated by the M+L mechanism, with very little contribution from the S-cones; however, we found that for all participants L-M also contributes negatively. Together these two findings suggest that the M-cones, more than L-cones, drive the detection of motion-luminance. Finally, although the S-cones contribute little to motion-luminance, which is in agreement with the literature on this point, nearly all participants

show significant S-cone contribution, with the polarity of the contribution either positive or negative for different participants. The relevance of this point, namely, the small but significant contribution with different polarity for different subjects, remains to be explored.

In our next chapter, we will consider how our method to derive highly precise measurements of equiluminant lights compares with another, widely used method, that of Heterochromatic Flicker Photometry.

Chapter 2

Comparing equiluminant settings derived using minimum motion versus heterochromatic flicker photometry

2.1 Introduction

In a typical application of Heterochromatic Flicker Photometry (HFP), the participant views a stimulus that alternates rapidly in time between two lights of different color; the participant then adjusts the intensity of one of the two lights (i.e., the amplitude of the light's spectrum) so as to minimize the sensation of flicker produced by the alternating lights. HFP has long been the standard psychophysical method for finding equiluminant lights (Ives, 1912a,b,c,d,e, 1914, 1917). Indeed, the luminous efficiency function $V(\lambda)$ developed by the CIE in the early part of the 20^{th} century was based largely on HFP.

Another commonly used method for obtaining equiluminant lights is the "minimally distinct border" (MDB) method (Boynton and Kaiser, 1968). In this task, the participant views

abutting patches of different colored lights and adjusts the intensity of one of the lights so as to render the border between the two lights minimally distinct. As shown by (Wagner and Boynton, 1972), equiluminant settings derived using HFP agree well with those derived using the MDB method in the sense that the two methods yield similar estimates of the luminous efficiency function.

To our knowledge, however, there has been no study comparing the equiluminance settings derived using HFP with those derived using the minimum motion method. This is the purpose of the current study. From each of three participants, we will obtain two data sets, one using the minimum motion method, the other using HFP. The procedure using the minimum motion method will be as described in Chapter 1. The procedure using HFP will be precisely analogous. Each data set will enable us to estimate 20 lights, each (i) maximally saturated on our display device and (ii) equiluminant (either motion-equiluminant or flicker-equiluminant) to the same fixed, gray standard light. As we did in Chapter 1, we will:

1. project each of these 20 lights into the 3-dimensional space spanned by the red, green and blue primaries of our display device,

2. find the plane that minimizes the sum of squared deviations of these 20 projections,

3. take the vector normal to this plane as an estimate of the (motion- or flicker-) luminance axis, and

4. compute the deviations of the 20 projections of the (motion- and flicker-) equiluminant lights from the best-fitting plane.

5. The central empirical questions are:

   (a) does the pattern of the deviations for the 20 flicker-equiluminant lights match the pattern for the motion-equiluminant lights?

(b) do the flicker-luminance axes of our three participants match their motion-luminance axes?

2.2 Methods

2.2.1 Participants

Three subjects participated in our experiments. All reported normal or corrected-to-normal vision. The UC Irvine Institutional Review Board approved the experimental procedures, and all participants gave signed consent.

2.2.2 Equipment

An iMac desktop computer running OS X version 10.6.8 with a 3.06 GHz Intel Core 2 Duo processor and 4 GB memory capacity was used for stimuli presentation and data collection. The computer was equipped with an ATI Radeon HD 4670 graphics chip. The monitor had a resolution of 1920 x 1080 and a viewable diagonal measure of 21.5 inches.

2.2.3 Stimuli

The stimulus used in the HFP task is illustrated in Fig. 2.1. The stimulus alternated between two movie frames (schematized in Fig. 2.1 A and B) at the rate of 30 frames/sec., yielding a flicker frequency of 15 Hz (matched to the temporal frequency of the motion stimulus used in the minimum motion task). Each half of the display alternated between the colored test light T (which is shown as green in Fig. 2.1 but which varied in hue across different experimental conditions) and the gray standard light G which had pixel value $v_G = (175, 175, 175)^T$ on

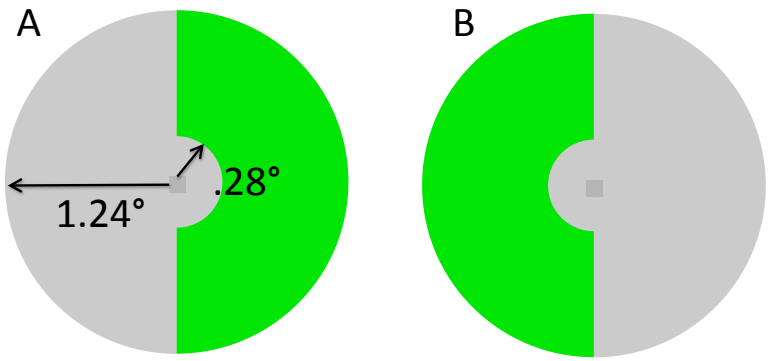

Figure 2.1: *The stimulus used in the HFP task.* The stimulus alternated at the rate of 15 Hz (30 frames per sec) between the stimuli shown in A and B. As indicated in A, the inner radius of each of the left and right half-annuli was 0.28° and the outer radius was 1.24°. The participant fixated the central cue spot and adjusted the color of the non-gray region by pressing either the right arrow or the left arrow to minimize the flickeriness of the stimulus.

our display device (photometric luminance $52cd/m^2$). However, the left and right sides of the display alternated in opposite phases; this created a vertical border down the middle of the stimulus that was present throughout the display. (We conjectured that introducing the spatial border into the stimulus might make the judgment easier; however, at the flicker rate of 15 Hz, the border was barely visible, even when the percept of flicker was strong, suggesting that adding this border had little effect.) The outer and inner diameters of the annular, flickering region subtended 1.24 and 0.28 deg. of visual angle respectively. A given stimulus presentation included 4 cycles of the temporally periodic stimulus.

To verify that the results were not due to a different spatial frequency of the two stimuli, we ran an additional control condition in which the stimuli had the exact same spatial frequency as the stimuli in the experiments described in chapter 1. The results were virtually the same, and are discussed in the Results section.

The stimulus used in the minimum motion task is illustrated in Fig. 1.3 of Chapter 1. To reiterate the key parts of the description from Chapter 1: The stimulus is periodic in time with four frames, each of which comprises an annular section of a radial square-wave grating that runs through four cycles per circumference. Frame 1 is shown on the left of Fig. 1.3.

In Frame 2, the square wave has been shifted one quarter cycle; if we think of this shift as clockwise (which corresponds to downward on the right side of Fig. 1), then the sections that were dark in Frame 1 have become green, and the sections that were light in Frame 1 have become gray. Frame 3 is identical to Frame 1, except that the dark and light bars switch roles, and Frame 4 is identical to Frame 2, except that the green and gray bars switch roles. On a given trial, the subject fixated a small cue spot in center of the screen and pressed the space bar to initiate a trial. Eight frames of the annular motion stimulus were then displayed at the rate of 60 frames per second, producing a stimulus temporal frequency of 15 Hz. The outer and inner diameters of the motion annulus were 1.24 and 0.28 deg. respectively. The stimulus regions colored medium gray in Fig. 1.3 were populated with standard light G. The stimulus regions colored light gray and dark gray in Fig. 1.3 were achromatic with Weber contrasts 0.06 and -0.06 respectively relative to G. The stimulus regions colored green in Fig. 1.3 were populated with a test light T that was varied from trial to trial.

2.2.4 Conditions

In each of the minimum motion and the HFP tasks the participant was tested in 20 different conditions, two conditions for each of the ten perimeters P_k, $k = 1, 2, \cdots, 10$ described in the Procedure section of Chapter 1. Each of the perimeters P_k comprises a cyclically ordered set of 400 pixel values constructed so as to intersect the set of lights (either flicker- or motion-) equiluminant to the standard light G at two points (each of which will be maximally saturated on our display device). (Fig. 1.4 shows the lights in all 10 perimeters P_k.) In one condition using the lights in P_k, data are collected to enable estimation of one of these two points; in the other condition, data are collected to enable estimation of the other.

In the minimum motion task, the participant first adjusted the color of a square subtending ≈ 0.5 deg. visual angle, placed in the center of the screen, to be roughly equiluminant with a

homogeneous background of standard light G. By repeatedly pressing the left-arrow (right-arrow) key, the participant could move around P_k in one (the other) direction. When the subject judged that the colored spot was nearly equiluminant to the background, he/she pressed the "o" key. The point of this initial adjustment procedure was to focus all of the trials in the motion-direction task (described below) in the neighborhood of the equiluminant light to be estimated. Following this initial, rough adjustment procedure, the participant then performed 60 trials in the minimum motion task. On the first of these trials, the test light T was assigned the final color to which the participant had adjusted the spot. On each trial, the participant fixated the small central cue spot and initiated a stimulus presentation with a button-press. The participant then viewed a single stimulus display and pressed the right arrow key (left arrow key) if the stimulus appeared to move clockwise (counterclockwise). No feedback was given. A 1-up-1-down staircase procedure was used to concentrate observations in the neighborhood of the light that was motion-equiluminant to G. Specifically, if the current trial used a test light T with pixel value (r, g, b), then if the participant's response indicated that T was higher (lower) in motion-luminance than G, the next test-light pixel value was the nearest one to (r, g, b) in the current perimeter lower (higher) in photometric luminance.

In the HFP task, in each of the 20 different color conditions, the participant performed 10 adjustments, each of which yielded a separate estimate of the particular point on a fixed side of one of the 10 perimeters that was flicker-equiluminant to the standard light G. To perform a given adjustment, the participant first (as at the start of a given condition in the motion task) adjusted the color of a square subtending ≈ 0.5 deg. of visual angle, placed in the center of the screen, to be roughly equiluminant with a homogeneous background of standard light G, pressing the "o" key when he/she judged the spot to be roughly equiluminant to the background.

After the participant pressed the "o," a small cue spot, slightly darker than the background, appeared in the middle of the screen. The participant could then press any key to initiate a single, 4-cycle, stimulus display. Following the stimulus presentation, the participant could press the right arrow key to increase the luminance of T or the left arrow key to decrease it. Pressing either key also initiated the next display. When the participant was satisfied that the current luminance level of T minimized the flickeriness of the stimulus, he/she pressed the "o" key to save the current setting.

2.3 Results

As in Chapter 1, we start by mapping the colors estimated to be flicker- and motion-equiluminant into "B-codes," where B is the basis of the space spanned by the spectra of the red, green and blue fundamentals of our display device. These fundamentals are shown in Fig. 1.1 of Chapter 1, and the three components of the orthonormal basis B used in the current experiments are plotted in Fig. 1.2 of Chapter 1. The B-code of a light with spectrum Q in the space spanned by B is the vector $q \in \mathbb{R}^3$ such that $Q = Bq$.

Fig. 2.2 shows the 20 B-codes corresponding to the flicker-equiluminant lights estimated for one participant p. The motion-equiluminant lights for this same participant are plotted in Fig. 1.5 of Chapter 1. The thick black bar shows the estimated flicker-luminance axis $\mathcal{L}_f(p)$. The thin black lines all lie in the estimated flicker-equiluminant plane, which is the plane of lights whose codes q satisfy $\mathcal{L}_f(p)^T q = \mathcal{L}_f(p)^T q_G$ (where q_G is the B-code of the standard light G with pixel value $v_G = (175, 175, 175)$). The colored circles show the codes of the individually estimated, flicker-equiluminant lights. Comparing Fig. 2.2 and Fig. 1.5 it is evident that the inclination of the flicker-luminance axis differs slightly from that of the motion-luminance axis for this participant. As was true in the case of motion-luminance, some of the flicker-equiluminant lights seem to deviate from the estimated flicker-

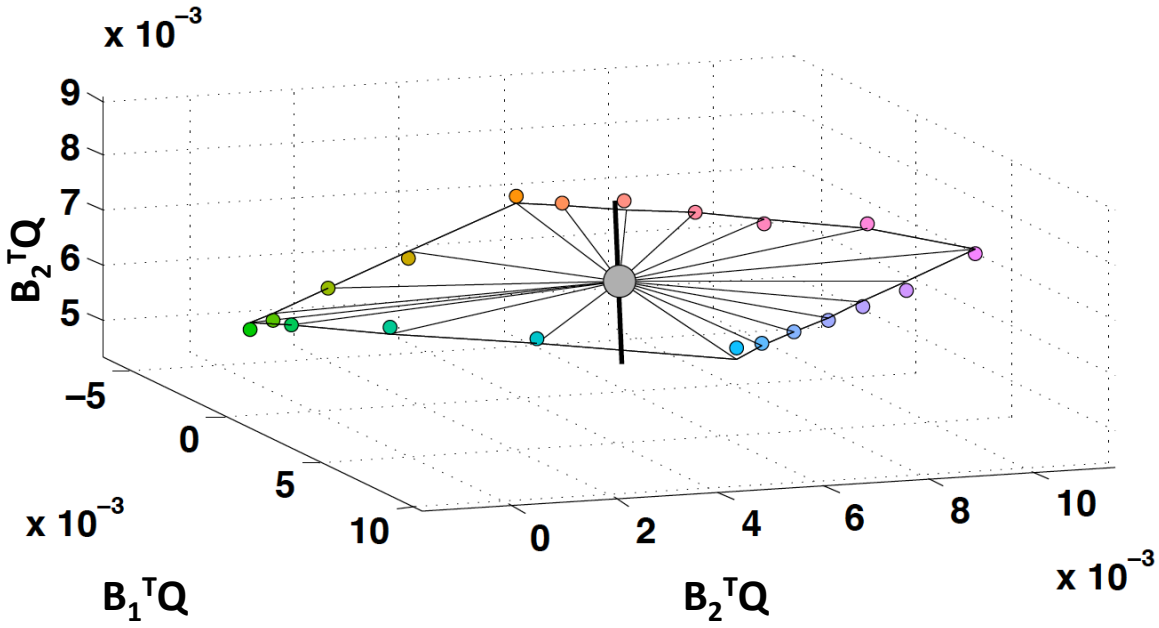

Figure 2.2: *The 20 B-codes corresponding to the flicker-equiluminant lights estimated for one participant p.* The thick black bar shows the estimated flicker-luminance axis $\mathcal{L}(p)$. The thin black lines all lie in the estimated flicker-equiluminant plane, which is the plane of lights whose codes q satisfy $\mathcal{L}(p)^T q = \mathcal{L}(p)^T q_G$ (where q_G is the B-code of the standard light G with pixel value $v_G = (175, 175, 175)$). The colored circles show the codes of the individually estimated, flicker-equiluminant lights.

equiluminant plane. However, it is difficult to tell from Fig. 2.2 whether these deviations are significant. It is also difficult to tell whether the estimate of the flicker-luminance axis differs significantly from that of the motion-luminance axis.

We first take a closer look at the deviations of the B-codes of both the flicker- and motion-equiluminant lights from the corresponding equiluminant planes. These deviations of the B-codes of the flicker-equiluminant (motion-equiluminant) lights are plotted for all three participants in the left (middle) three panels of Fig. 2.3. To facilitate comparison, the right

panel plots the results from the first panel (black line joining triangles) and the second panel (gray line joining circles) together, omitting error bars.

There are several things to note about Fig. 2.3. First, in most cases the 95% Bayesian credible intervals are smaller for the motion-equiluminant B-codes than are the 95% confidence intervals for the flicker-equiluminant B-codes. This is striking because typically it takes less time (approximately 2 minutes) to collect the data in a given condition with the minimum motion task than it does to collect the data in a condition with the HFP task (times varied, roughtly 5 minutes, but more for some subjects, for some conditions, markedly more). Second, the deviations of the motion-equiluminant B-codes from the estimated motion-equiluminant plane tend to be more dramatic with more clearly defined structure than do the deviations of the flicker-equiluminant B-codes from the estimated flicker-equiluminant plane.

In the current chapter, we write $\mathcal{L}_{f,p}$ and $\mathcal{L}_{m,p}$ for the flicker- and motion-luminance axes for participant p. Properly speaking, the term "motion-luminance axis" (the same goes for "flicker-luminance axis") would only make sense if the B-codes of all of the lights that are motion-equiluminant to a given light reside in a plane. This is clearly not true for either motion- or flicker-equiluminance, though more dramatically in the case of motion-luminance than flicker-luminance. Nonetheless, as suggested by Fig. 2.2 as well as by Fig. 1.5 of Chapter 1, the best-fitting plane accounts for a high proportion of the variance in the B-codes of the 20 equiluminant lights for both the minimum motion and the HFP tasks. We thus write $\mathcal{L}_{f,p}$ to refer to the function within the space spanned by the human cone fundamentals for which the plane of lights orthogonal to $\mathcal{L}_{f,p}$ and containing G provides the best planar description (in a least-squares sense) of the set of lights set to be flicker-equiluminant to G. Similarly, we write $\mathcal{L}_{m,p}$ to refer to the analogous function for which the plane of lights orthogonal to $\mathcal{L}_{m,p}$ and containing G provides the best description of the set of lights estimated to be motion-equiluminant to G.

Strikingly, for all participants, the adjacent B-codes (of the bluish lights), plotted at hue angles 0 and 38 deviate positively from the estimated equiluminant plane; this is true both for flicker and even more dramatically for motion. In the case of motion-luminance, for example, this means that lights with these particular hues must include greater levels of $\mathcal{L}_{m,p}$ than lights with other hues in order to be made motion-equiluminant to G. Especially striking is the difference in motion-luminance between the two hues plotted at hue angles 0 and 38 vs. the next three lights plotted at angles 44, 55 and 66. Although these three lights seem similar in hue to the first two, they deviate strongly and negatively in motion-luminance from the estimated motion-equiluminant plane. In the flicker-equiluminant settings, we observe a similar decrease in $\mathcal{L}_{f,p}$ in the B-codes of the lights ranging from blue to lavender. However, this decrease is not as steep as it is for the motion-equiluminant lights. This pattern is articulated most cleanly in the results for participant P1 but is evident for all three participants.

B is a basis of the space spanned by the red, green and blue primaries of the particular monitor we used for the current experiments. There is no reason to suppose that this space is identical to the space spanned by the human cone fundamentals. Therefore, because we assume that the motion-luminance axis for a given participant resides in the space spanned by the human cone fundamentals, we cannot use B-codes to represent the motion-luminance axes of our participants. As we did in Ch. 1, we will use the basis F derived from the Stockman-Sharpe 2 deg. cone fundamentals to characterize the functions $\mathcal{L}_{m,p}$ and $\mathcal{L}_{f,p}$ estimated for each of our three participants p. The three orthonormal basis functions F_{M+L}, F_S and F_{L-M} are defined precisely in Ch. 1. We write F for the matrix whose three columns (from left to right) are F_{M+L}, F_S,and F_{L-M}. For any function Q in the space spanned by the human cone fundamentals, we call the column vector F^TQ the F-code of Q.

It is easy to check that F_{M+L}, F_S and F_{L-M} compose an orthonormal basis of the space spanned by the Stockman-Sharpe cone fundamentals. Let F be the matrix whose three

columns are the functions F_{M+L}, F_S,and F_{L-M}. For any function Q in the space spanned by the human cone fundamentals, we call the column vector F^TQ the F-code of Q.

Fig. 2.4 plots the F-codes of the motion-luminance and flicker-luminance axes for all 3 participants. In panel A, all three dimensions of the F-code are shown. The blue (red) crosses give the motion-luminance axis F-codes (flicker-luminance axis F-codes) for participants P1, P2 and P3. The red crosses give the flicker-luminance axis F-codes for participants P1, P2 and P3. Error bars (signaled by the arms of the crosses) are 95% Bayesian credible intervals aligned with the principal components of posterior density. The error bars in the third dimension (orthogonal to the crosses) are too small to be visible. Note that for all participants both the flicker-luminance axis and the motion-luminance axis are dominated by F_{M+L}: $F_{M+L}^T\mathcal{L}_{m,p}$ and $F_{M+L}^T\mathcal{L}_{f,p}$ are around 0.95 for all participants; by contrast, $F_S^T\mathcal{L}_{m,p}$ and $F_S^T\mathcal{L}_{f,p}$ are near 0, $F_{L-M}^T\mathcal{L}_{m,p}$ is around -0.07, and $F_{L-M}^T\mathcal{L}_{f,p}$ is around -0.13.

Strikingly, the F-code of the flicker-luminance axis differs significantly from the F-code of the motion-luminance axis for each of our three participants. In each case the credible intervals for the two axes are far removed from each other. It is also suggestive that the F-code of the flicker-luminance axis of each of our three participants is displaced in roughly the same direction from the F-code of his/her motion luminance axis. For each participant p, the F-code of p's flicker-luminance axis (red) undergoes a slight negative shift in $F_S^T\mathcal{L}_{f,p}$ and a larger negative shift in $F_{L-M}^T\mathcal{L}_{m,p}$ from p's motion-luminance axis (blue). When we ran the control condition using stimuli with the same spatial frequency as the stimuli used in the minumum motion experiment described in Chapter 1, we found that the results were the same except that that the value of $F_S^T\mathcal{L}_p$ goes from around -0.03 for the original data to 0.03 for the data from the control experiment. Indeed, the $F_{L-M}^T\mathcal{L}_{f,p}$ values for each of our three participants are all more negative than any of the nine $F_{L-M}^T\mathcal{L}_{m,p}$ values observed in Ch. 1, Fig. 1.9.

2.4 Discussion for Chapter 2

When comparing two different performance-based methods to estimate luminance axis, we discover important differences in both the performance-based luminance axis and the performance-equiluminant lights they yield. In Chapter 1 we show that motion-equiluminant lights do not lie on a plane, and we find as much for flicker-equiluminant lights. One important consequence is that, contrary to previous assumptions about the relationship between performance-luminance and performance-equiluminant lights, a light with spectrum Q cannot be perfectly modeled by the inner product of a function such as V_λ with Q. In the case of flicker-equiluminant lights, however, the deviations from the planar fit are small; in this case, the approximation provided by the inner product is likely to be pretty good provided one uses the right function. For all three subjects, we observe a similar pattern of deviation of the estimated flicker-equiluminant lights from the best-fitting plane. Notably, some of the blue hues, hue angles 0 and 38, are above the plane, some of the lilacs, adjacent to them, hue angle 44, are distinctly below, and then there is a gradual "bounce" back for the rest of the lights. The undulating pattern distinct on the motion hues is present, although much less pronounced. There are marked differences between the set of hues obtained via motion and the set of hues obtained via flicker. For hues obtained via flicker, the confidence intervals are in some cases quite large, compared with the very small Bayesian credible intervals for hues obtained via motion. For both methods, the hue centered at angle 0 is firmly above the best-fitting plane; however, for the hues obtained via motion, the "drop" is more pronounced and clearly defined for all subjects. One might speculate that the difference between flicker-equiluminance might be attributable to difference between neurons that gauge "flickeriness" from the ones that detect first-order motion. A priori, there is no compelling reason why it should be the case that the sensitivity to light-spectral variations of the neurons used to sense "flickeriness" must be matched to the sensitivity of the neurons used to sense first-order motion. Indeed, our results show a shift in the flicker-luminance axis relative to the motion

luminance axis for all three participants; such shift suggests that the neurons recruited for judging flickeriness have higher relative sensitivity to M-cone activations vs L-cone activations than do the neurons used to judge first-order motion direction. Our results clearly show deviations of the B-codes of our equiluminant lights from the best-fitting plane. Therefore, the flicker-luminance of a light with spectrum Q cannot be perfectly modeled by the inner product of a function such as V_λ with Q. The situation is the same for motion-luminance, as discussed in Chapter 1. We don't know, at the moment, what kind of model could account for it, but our results impose important constrains. The model should account for extremely abrupt changes in deviations away from the best-fitting plane for very small changes in hue, in some cases of just one step in our hue directions. It also should account for both positive and negative deviations. Finally, we note that whenever a particular hue shows a positive deviation from the best-fitting plane, for both motion- and flicker-luminance, this means that in order for this hue to be equiluminant to the standard gray light G, it needs to project higher on the motion-luminance axis than other hues. This means that a light of this hue projected onto the plane would actually be lower in motion-luminance than other hues. Stated succinctly: a positive deviation means a particular hue tends to be lower in motion-luminance than other hues, and a negative deviation means a particular hue tends to be higher in motion-luminance than other hues.

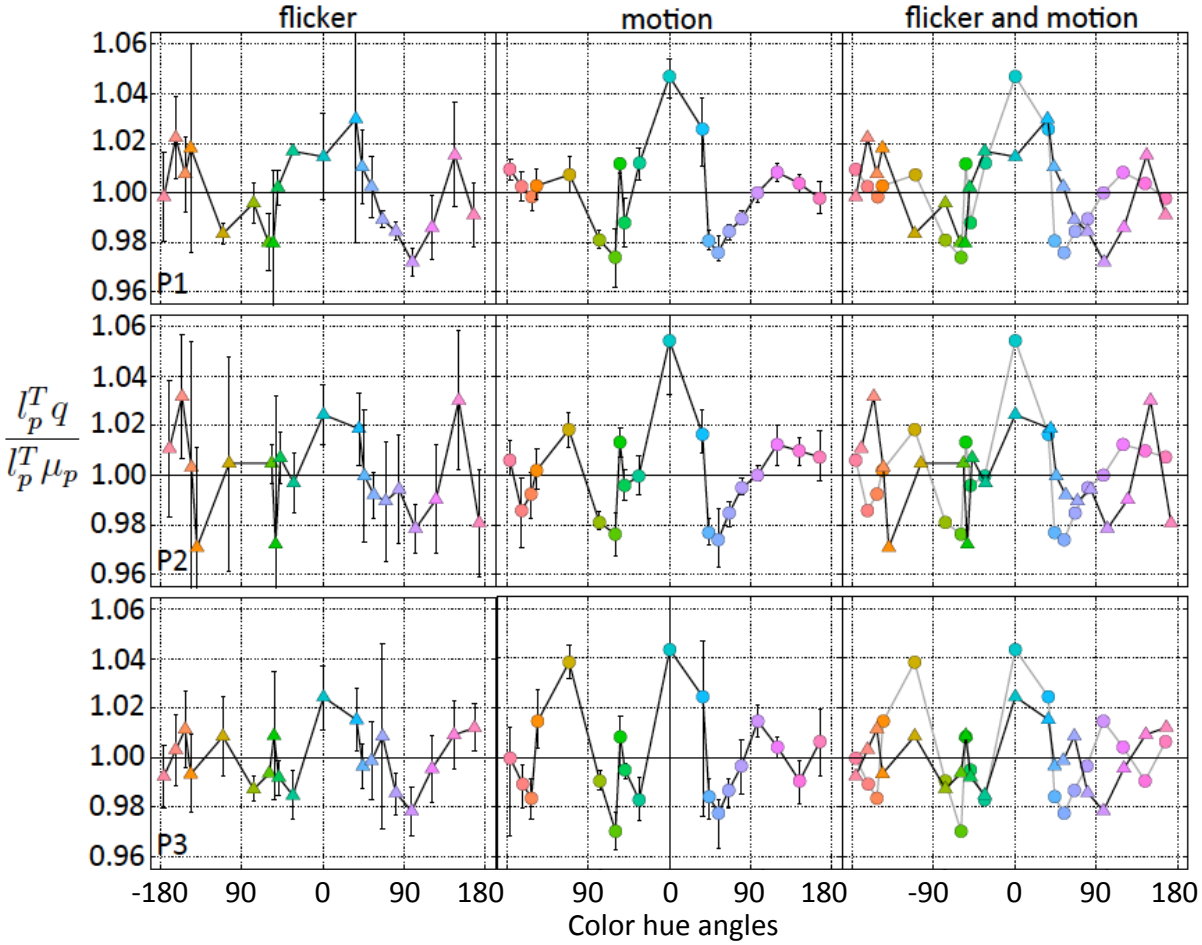

Figure 2.3: *The mean deviations of the 20 B-codes corresponding to the flicker- and motion-equiluminant lights.* Each row of three panels corresponds to a different participant. The leftmost panel shows the deviations of the B-codes of the 20 flicker-equiluminant lights from the estimated flicker-equiluminant plane. Error bars in this figure are 95% confidence intervals (derived from a t distribution with 9 degrees of freedom) for the mean of the 10 settings produced by the participant. The middle panel shows the deviations of the B-codes of the 20 motion-equiluminant lights from the estimated motion-equiluminant plane. Error bars in this figure are 95% Bayesian credible intervals (derived using Markov chain Monte Carlo simulation) for the mean of cumulative normal psychometric function used to fit the psychometric data from a given condition. To facilitate comparison, the right panel plots the results from the first panel (black line joining triangles) and the second panel (gray line joining circles) together, omitting error bars.

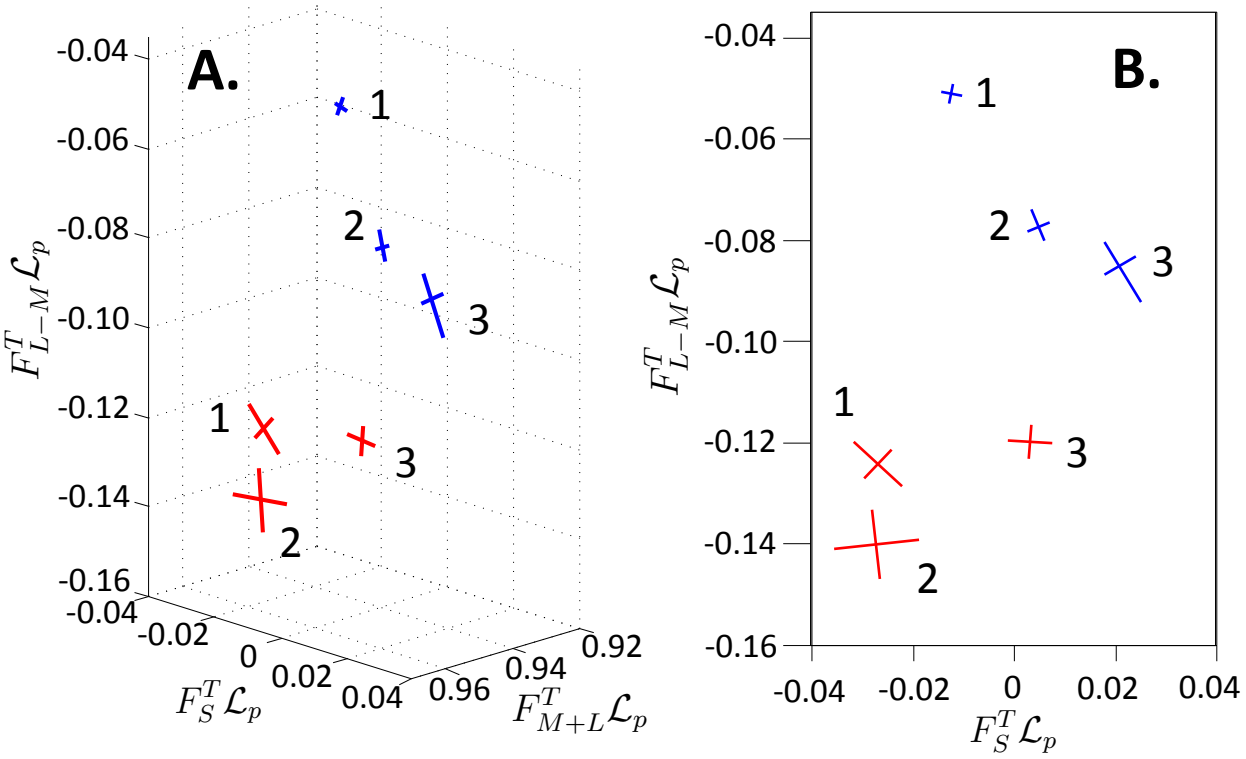

Figure 2.4: *The F-codes of the motion-luminance and flicker-luminance axes for all 3 participants.* A. All three dimensions are shown. The blue (red) crosses give the motion-luminance axis F-codes (flicker-luminance axis F-codes) for participants P1, P2 and P3. Error bars (signaled by the arms of the crosses) are 95% Bayesian credible intervals aligned with the principal components of posterior density. Note that for all participants both the flicker-luminance axis and the motion-luminance axis are dominated by F_{M+L}. B. The projection of the motion-luminance axis F-codes (blue) and flicker-luminance axis F-codes (red) of participants P1, P2 and P3 into the F_S, F_{L-S} plane.

Chapter 3

Preattentive mechanisms for colored white noise textures

3.1 Introduction

What color-selective mechanisms exist in human vision? This question has been the focus of a great deal of research effort, and yet no clear consensus has emerged. Substantial evidence supports the theory, originally proposed by Krauskopf et al. (1982) that human vision comprises three basic post-retinal mechanisms (Krauskopf et al., 1986; Krauskopf and Gegenfurtner, 1992; Krauskopf et al., 1996a,b; Gegenfurtner and Kiper, 1992; Webster and Mollon, 1991, 1994). It was proposed that these three mechanisms combine the responses of the short, medium and long-wavelength cone classes, S, M and L additively as follows:

1. The *luminance* mechanism sums the responses of the M and L cones: $M + L$.

2. The *blue-yellow* mechanism takes the difference between the S-cone response and the summed M and L cone responses: $S - (M + L)$.

3. The *red-green* mechanism takes the difference between the L- vs. M-cone responses: $L - M$.

The weights with which the different cone-class responses were combined by these three mechanisms were not well-constrained by the data available to Krauskopf et al. (1982). The important point was that these mechanisms combine cone-responses linearly.

Eskew (2009) emphasizes, however, that it really does not make sense to think of the red-green mechanism (a similar argument applies to the yellow-blue mechanism) as a single mechanism; rather we should think about it (at least for purposes of psychophysical analysis) as two, complementary, half-wave rectified mechanisms:

1. A *red half-axis* mechanism whose response is given by

$$[L-M]^{+} = \begin{cases} L-M & \text{if } L-M > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{3.1}$$

2. And a *green half-axis* mechanism whose response is given by

$$[L-M]^{-} = \begin{cases} |L-M| & \text{if } L-M < 0 \\ 0 & \text{otherwise.} \end{cases} \tag{3.2}$$

The argument supporting this contention involves two assumptions that are intrinsic to the idea of a mechanism: the *univariance* assumption and the *labeled line* assumption:

1. The univariance assumption holds that the output produced by a mechanism in response to a light is a single number: the activation produced in the mechanism by the light.

2. The labeled line assumption (Graham, 1989; Watson and Robson, 1981) holds that two lights will appear different if and only if they produce significantly different levels of activation in some mechanism.

Suppose, for example, that B is the spectrum of a background light. Let δ_1 and δ_2 be functions mapping the visible wavelengths into $\mathbb{R}$, and consider a stimulus consisting of a dot with spectrum $B+\alpha_1\delta_1$ next to a dot with spectrum $B+\alpha_2\delta_2$ on an otherwise homogeneous background with spectrum B. Suppose all mechanisms except one, that we shall call M, are blind to variations in the amplitude of each of δ_1 and δ_2. In this case, for any values of α_1 and α_2, the response of any mechanism other than M to the lights $B+\alpha_1\delta_1$ and $B+\alpha_2\delta_2$ is equal to its response to B. Suppose, however, that M is sensitive (in different degrees) to both δ_1 and δ_2. For concreteness, suppose M is more sensitive to variations in the amplitude of δ_2. The univariance assumption requires that for any α_1 there exists a value of α_2 such that the spots with spectra $B+\alpha_1\delta_1$ and $B+\alpha_2\delta_2$ appear identical. In particular, this will be true for α_2 adjusted so that the two spots produce equal activation in M. On the other hand, the labeled line assumption implies that if there exists a mechanism, M' distinct from M that is sensitive to variations in the amplitude of δ_1 but not to variations in the amplitude of δ_2, then no matter how one adjusts the amplitude α_2, there will always be a perceptual difference between the lights with spectra $B+\alpha_1\delta_1$ and $B+\alpha_2\delta_2$ because these two lights produce different levels of activation in M'.

Eskew (2009) observes that if the background light with spectrum B is gray, and the dots with spectrums $B+\alpha_1\delta_1$ and $B+\alpha_2\delta_2$ are red and green respectively, then even if the amplitudes α_1 and α_2 are both at detection threshold (under additional conditions designed to insure that performance is mediated by a single mechanism), the two dots appear different (Eskew et al., 2001; Krauskopf et al., 1986; Mullen and Kulikowski, 1990). This leads Eskew to conclude that human vision should be seen as possessing distinct mechanisms activated by δ_1 and δ_2. He refers to such mechanisms as linear, half-axis mechanisms.

On the other hand, suppose that the $L - M$ mechanism hypothesized by Krauskopf et al. (1982) can produce negative activations. This might be accomplished in several ways. First, the mechanism might use neurons that maintain a non-zero resting firing-rate. A response of 0 would then signal that the light impinging on the retina at a given point in space has $L-M$ activation equal to that of the background to which the participant is adapted. Negative activations (indicating that the light deviates in the green direction from the adapting level) would be signaled by deviations in firing below the resting state firing level, and positive activations (indicating that the light deviates in the red direction from the adapting level) would be signaled by deviations above the resting state firing level. The problem with a neuronal architecture of this sort is that it is costly: substantial resources are required to sustain a non-zero resting-state firing rate in a large population of neurons. A different strategy would make use of separate red and green half-axis channels precisely of the sort suggested by Eskew. Under this strategy, the red half-axis channel would transmit the positive deviations of the $L - M$ mechanism away from the adapted level, and the green half-axis channel would transmit the negative deviations.

Let us suppose that human vision does possess separate red and green half-axis mechanisms. Under one scenario, top-down attention might have entirely separate access to these two mechanisms for purposes of making judgments about various images. It certainly feels, for example, as though one can selectively attend to the greens in a painting without being influenced by the reds, and vice versa. This experience suggests separate access to mechanisms selective for red and for green. Another possibility, however, is that our hypothetical red and green half-axis mechanisms are bound together in a way that precludes separate access, with the red half-axis mechanism coding the positive activations in a single opponent mechanism and the green half-axis mechanism coding the negative activations.

Several previous studies provide support for the idea that participants have separate access to red and green half-axis mechanisms. First, adapting to temporal modulations of

chromaticity that follow a sawtooth profile (e.g., repeatedly changing gradually to red, then changing abruptly to green and changing gradually again to red, etc.) can increase thresholds differentially for the color receiving the abrupt onset vs. the gradually changing color (Krauskopf et al., 1986; Krauskopf and Zaidi, 1986). It has also been shown that briefly flashed reddish targets are much more effectively masked by reddish than by greenish noise, and vice versa (Sankeralli and Mullen, 2001). In addition, it has been shown that targets that deviate from gray in the red direction vs. those that deviate in he green direction are detected by different mechanisms at threshold in the fovea (Gowdy et al., 1999a,b).

The current study uses a different experimental paradigm involving discrimination of chromatic textures to seek evidence of red- vs. green-selective half-axis mechanisms. This paradigm was originally used by Silva and Chubb (2014) to analyze the mechanisms sensitive to textures composed of different grayscales.

3.2 Methods

3.2.1 Participants

A total of 9 subjects (4 female) participated in our experiments. All reported normal or corrected-to-normal vision. That none had color deficiencies was verified using the minimum motion task (see below, section 2.3). The UC Irvine Institutional Review Board approved the experimental procedures, and all participants gave signed consent.

3.2.2 Equipment

An iMac desktop computer running OS X version 10.6.8 with a 3.06 GHz Intel Core 2 Duo processor and 4 GB memory capacity was used for stimuli presentation and data collection.

The computer was equipped with an ATI Radeon HD 4670 graphics chip. The monitor had a resolution of 1920 x 1080 and a viewable diagonal measure of 21.5 inches.

3.2.3 Color Palette acquisition

For each participant, we used a minimum motion task to acquire a palette of colors motion-equiluminant with the background. The background (with respect to which all our lights were motion-equiluminant) had a photometric luminance of $52cd/m^2$, as measured with a spectrophotometer PR-670. This process is described in detail in Chapter 1. Importantly, each participant had lights (colors) that were both unique to the participant, and motion-equiluminant to the background. For the present experiment, we used a set of motion-equiluminant colors that isolated the L-M cardinal axis of DKL space. Specifically, the lights in this set

1. all produced equal activation in the (Stockman-Sharpe, 2 deg.) S-cone fundamental.
2. were all motion-equiluminant to the achromatic light with photometric luminance of $52cd/m^2$.

Thus, these lights were chosen to produce differential activation only in whatever mechanisms exist in human vision sensitive to variations in the difference between L-cone vs. M-cone activations. Eight lights were used for each participant. These lights projected onto locations -3.5α, -2.5α, -1.5α, -0.5α, 0.5α, 1.5α, 2.5α, 3.5α, of the $L-M$ cardinal axis of DKL space, where α was chosen to be as large as possible on our display device given that location 0 corresponded to the achromatic light with photometric luminance of $52cd/m^2$. We will write Ω for the set containing these 8 colors.

3.2.4 Ω-scrambles

Stimuli were composed of chromatic textures called Ω-scrambles. An Ω-scramble is a patch of texture comprising small square elements, each of which is colored with one of the colors in Ω. The "histogram" of the scramble is the probability distribution that gives the proportions with which the different colors in Ω appear in the patch. In practice, to generate a patch of Ω-scramble comprising N spatial elements with histogram as close as possible to p, one loads a virtual urn with N copies of colors drawn from Ω in proportions conforming as closely as possible to p. These colors are then assigned randomly without replacement to the N locations of the patch.

We will write U for the uniform histogram (i.e., $U(\omega) = \frac{1}{8}$ for all $\omega \in \Omega$).

In addition, we call any function $\rho : \Omega \to \mathbb{R}$ a perturbation if

1. $\sum_{\omega\in\Omega} \rho(\omega) = 0$, and
2. $|\rho(\omega)| < \frac{1}{8}$ for all $\omega \in \Omega$.

These two conditions insure that $U + \rho$ and $U - \rho$ will both be probability distributions on Ω. If in fact $|\rho(\omega)| = \frac{1}{8}$ for some $\omega \in \Omega$, we call ρ a maximal perturbation.

We will assume that any visual mechanism differentially sensitive to Ω-scrambles can be characterized by a sensitivity function $f(\omega)$ that reflects the activation produced in the mechanism by a texture element of color ω. Under this assumption, the space-average activation produced in the mechanism by an Ω-scramble with histogram p is equal to

$$f \bullet p = \sum_{\omega\in\Omega} f(\omega)p(\omega), \tag{3.3}$$

and the difference in activation produced in the mechanism by scrambles with histograms p and q is $f \bullet (p - q)$. Thus, if $p = U + \rho$ and $q = U - \rho$, this difference in activation is $2f \bullet \rho$.

Our goal in the current experiment is to determine the sensitivity functions characterizing all of the mechanisms in human vision that are differentially sensitive to Ω-scrambles.

3.2.5 Stimuli

Example stimuli are shown in Figs. 3.1 and 3.2. Although no chin rest was used, subjects remained seated straight and with the head level, at a distance of approximately 95 cm. At this distance the outer diameter of the stimulus annulus subtended 6.8° of visual angle, and each texture element subtended 0.11° of visual angle.

3.2.6 What happened on a given trial

Each trial was initiated with a button press. There followed a 500 ms. blank gray screen with central cue spot slightly brighter than the background. The stimulus then appeared for 300 ms and was then replaced by the blank screen with the cue spot. During the stimulus presentation, the central cue spot remained visible in the blank region at the center of the annular texture patch comprising the stimulus. The Target disk could appear at one of eight fixed locations, equidistant from the cue spot: horizontally to the right or left of the cue spot, vertically above or below the cue spot, or at one of the four locations diagonally intervening between these four cardinal locations. After the display, the participant used the number pad keys to indicate the location of the target disk. The mapping was: “7” for up-left, “8” for up, “9” for up-right, “6” for right, “3” for down-right, “2” for down, “1” for down-left, “4” for left. A beep sounded after any incorrect response.

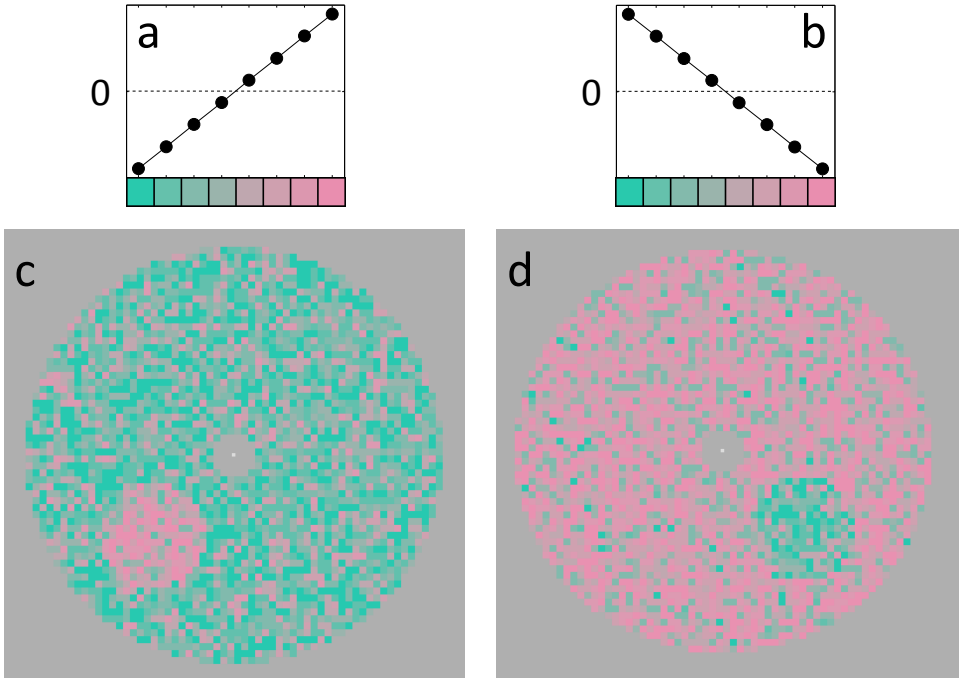

Figure 3.1: *Examples of stimuli from complementary conditions 1 and 2.* The perturbations (a) λ_1, and (b) $-\lambda_1$. The example stimulus in (c) has a target disk that is composed of Ω-scramble with histogram $U + A\lambda_1$, and a background annulus with histogram $U - A\lambda_1$, where the histogram amplitude A is chosen to make the perturbation $A\lambda_1$ maximal. The task is to indicate the location (using the keys in the number pad of the target disk). The roles of two different types of scramble have been reversed in the example stimulus in (d). In condition 1 (condition 2), each stimulus has a target disk with histogram $U + \rho$ and background annulus with histogram $U - \rho$ for some perturbation ρ that correlates strongly with λ_1 ($-\lambda_1$). In order to clearly indicate the nature of the search task in which the participant was engaged, the differences between the targets and backgrounds shown here are much stronger than they were in the actual experiment.

3.2.7 Experimental conditions

Each participant performed 2600 trials in each of four, separately blocked conditions. Each of these conditions was an individual application of the "seed expansion" method Chubb et al. (2012). The next section gives a brief overview of the method as it was used in one of the four conditions in the current experiment.

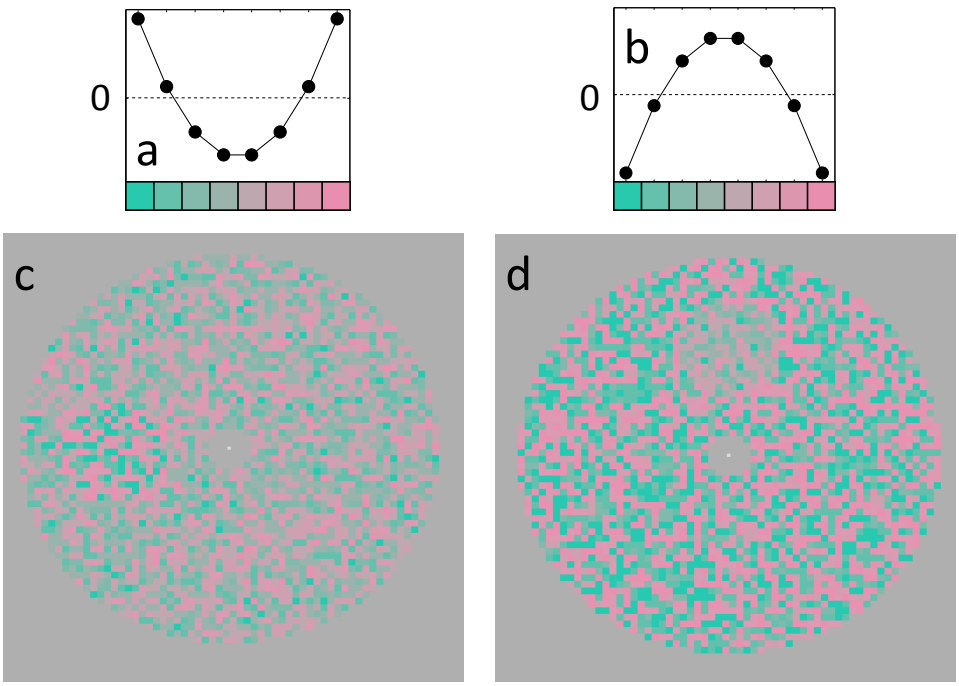

Figure 3.2: *Examples of stimuli from complementary conditions 3 and 4.* The perturbations (a) λ_2, and (b) $-\lambda_2$. The example stimulus in (c) has a target disk that is composed of Ω-scramble with histogram $U + A\lambda_2$, and a background annulus with histogram $U - A\lambda_2$, where the histogram amplitude A is chosen to make the perturbation $A\lambda_2$ maximal. The task is to indicate the location (using the keys in the number pad of the target disk). The roles of two different types of scramble have been reversed in the example stimulus in (d). In condition 3 (condition 4), each stimulus has a target disk with histogram $U + \rho$ and background annulus with histogram $U - \rho$ for some perturbation ρ that correlates strongly with λ_2 ($-\lambda_2$). In order to clearly indicate the nature of the search task in which the participant was engaged, the differences between the targets and backgrounds shown here are much stronger than they were in the actual experiment.

3.2.7.1 The use of the seed expansion method in the current experiment

In a given one of our four separately blocked conditions, a single dominant perturbation ϕ is used to produce the difference between the target disk vs the background annulus on each trial. We call ϕ the *seed* perturbation for the condition. (In conditions 1, 2, 3 and 4, the seed perturbations are respectively λ_1, $-\lambda_1$, λ_2 and $-\lambda_2$.) On any given trial in the condition with seed ϕ, the Ω-scramble filling the target disk will have a histogram $U + \rho$ for some perturbation ρ correlated strongly and positively with ϕ (all correlations are 0.894), and the annular background will have histogram $U - \rho$. Thus, the qualitative difference

between the target-disk vs the background will be similar from trial to trial. (In condition 1, as suggested by Fig. 3.1, the target will tend to be pinker than the background on every trial; in condition 2, the target will tend to be greener than the background on every trial; in condition 3, the target will tend to be higher in red-green variance ("contrastiness") than the background, and in condition 4, the target will tend to be lower in red-green variance than the background.) The point of separately blocking these different conditions is to encourage the participant to use top-down attention to optimize his/her strategy for the particular target-vs-background texture difference in each given condition. In particular, the model we will fit proposes that the participant combines information from his/her mechanisms to produce an "Ω-filter" f_ϕ that gives high values to colors $\omega \in \Omega$ that occur with high density in the target and low values to colors ω that occur with high density in the background. We assume that the participant applies f_ϕ to the stimulus on a given trial to produce a neural map in which the target location is highly activated in comparison to the background. We call the function f_ϕ the expansion of the seed ϕ.

We estimate the expansion f_ϕ using a general linear model. The regression variables are the values $\rho(\omega)$, for all $\omega \in \Omega$, and the linking function is a Weibull function. The specific assumptions are:

1. The salience of the target in the condition with seed ϕ on a trial in which the target has histogram $U + \rho$ (and the background has histogram $U - \rho$) is

$$\mathrm{Sal}_\phi(\rho) = f_\phi \bullet \rho. \tag{3.4}$$

2. The probability of a correct response on such a trial is $\Psi_\phi(\mathrm{Sal}_\phi(\rho))$, for the Ψ_ϕ the Weibull function defined by:

$$\Psi_\phi(x) = 0.125 + 0.855\left(1 - \exp\left(-x^{\beta_\phi}\right)\right). \tag{3.5}$$

About Ψ_ϕ, notice that

1. Chance performance is $\Psi_\phi(0) = 0.125$ because the task requires an 8-option forced choice.

2. Ψ_ϕ asymptotes at 0.98 instead of 1.0 to cover the possibility of "finger errors," i.e., errors participants make make even though they clearly see correct response.

3. Typically, one expects a Weibull function to have two free parameters. However, there is only one (β_ϕ) in Eq. 3.5. This is because the other typical parameter (a scalar used to divide x) can be absorbed into the expansion f_ϕ in Eq. 3.4.

3.2.7.2 The four seed conditions

To describe the perturbations used in these experiments, we identify the 8 colors $\omega \in \Omega$ ranging from red to green with the values $v_1, v_2, \cdots, v_8$ equal to $-1, -\frac{5}{7}, -\frac{3}{7}, -\frac{1}{7}, \frac{1}{7}, \frac{3}{7}, \frac{5}{7}, 1$. The Legendre polynomials are derived by applying Gram-Schmidt orthonormalization to the sequence of monomials $h_j(v) = v^j$, $j = 0, 1, \cdots, 7$. The Legendre polynomials of order $1, 2, \cdots, 7$ are listed in table 3.1.

k	$\lambda_k(v_1)$	$\lambda_k(v_2)$	$\lambda_k(v_3)$	$\lambda_k(v_4)$	$\lambda_k(v_5)$	$\lambda_k(v_6)$	$\lambda_k(v_7)$	$\lambda_k(v_8)$
1	-0.5401	-0.3858	-0.2315	-0.0772	0.0772	0.2315	0.3858	0.5401
2	0.5401	0.0772	-0.2315	-0.3858	-0.3858	-0.2315	0.0772	0.5401
3	-0.4308	0.3077	0.4308	0.1846	-0.1846	-0.4308	-0.3077	0.4308
4	0.2820	-0.5238	-0.1209	0.3626	0.3626	-0.1209	-0.5238	0.2820
5	-0.1498	0.4922	-0.3638	-0.3210	0.3210	0.3638	-0.4922	0.1498
6	0.0615	-0.3077	0.5539	-0.3077	-0.3077	0.5539	-0.3077	0.0615
7	-0.0171	0.1195	-0.3585	0.5974	-0.5974	0.3585	-0.1195	0.0171

Table 3.1: The Legendre polynomials of order 1 to 7.

The experiment comprised 4 different conditions whose seed perturbations were $\phi = \lambda_1$, $-\lambda_1$, λ_2, $-\lambda_2$. (We had originally planned to also include conditions with seed perturbations

λ_3, and $-\lambda_3$. However, these conditions proved too difficult.) Examples of stimuli from the conditions with $\phi = \pm\lambda_1$ are shown in Fig. 3.1. Fig. 3.2 gives examples of $\phi = \pm\lambda_2$. (These example stimuli have the maximum possible histogram difference to make the texture differences characterizing the different conditions as obvious as possible.)

3.2.7.3 Trial-by-trial perturbations within a given seed condition

In each of the four conditions, the participant performed 2600 trials, 200 in each of 13 interleaved staircases. This section describes these staircases. Let $b_1 = \phi$, and let

$$b_2 = \begin{cases} \lambda_2 & \text{if } \phi = \pm\lambda_1 \\ \lambda_1 & \text{otherwise,} \end{cases} \tag{3.6}$$

and for $k = 3, 4, \cdots, 7$, let $b_k = \lambda_k$. Then we construct the perturbations

$$\eta_k^+ = \frac{b_1 + \frac{1}{2}b_k}{\|b_1 + \frac{1}{2}b_k\|} \qquad \text{and} \qquad \eta_k^- = \frac{b_1 - \frac{1}{2}b_k}{\|b_1 - \frac{1}{2}b_k\|} \tag{3.7}$$

for $k = 2, 3, \cdots, 7$. Note that each of the perturbations $\rho = b_1$, as well as $\rho = \eta_k^+$ and $\rho = \eta_k^-$ for $k = 2, 3, \cdots, 8$ is normalized. Note also that the correlation between ϕ and each of η_k^+ and η_k^- is 0.8944.

For each of the 13 perturbations $\rho = b_1, \eta_k^+, \eta_k^-$, $k = 2, 3, \cdots, 7$, psychometric data testing performance at localizing a target patch of $(U + A\rho)$-scramble in an annular background of $(U - A\rho)$-scramble was collected for various amplitudes A. Specifically, the staircase for a given perturbation ρ could visit the 30 histogram amplitudes $A = \frac{A_{max}}{30}, \frac{2A_{max}}{30}, \cdots, A_{max}$, for A_{max} the scalar for which the maximum absolute value of $A_{max}\rho$ is equal to $\frac{1}{8}$. Each staircase started at amplitude $A = \frac{A_{max}}{2}$ and ran for 200 trials. In each staircase, A was decremented whenever the previous two trials both yielded correct responses; otherwise A

was incremented. These 13 staircases were randomly interleaved to collect the 2600 trials of data in the condition with seed ϕ.

3.3 Results

The expansions f_ϕ achieved by all four participants in the different seed conditions are plotted in Fig. 3.3.

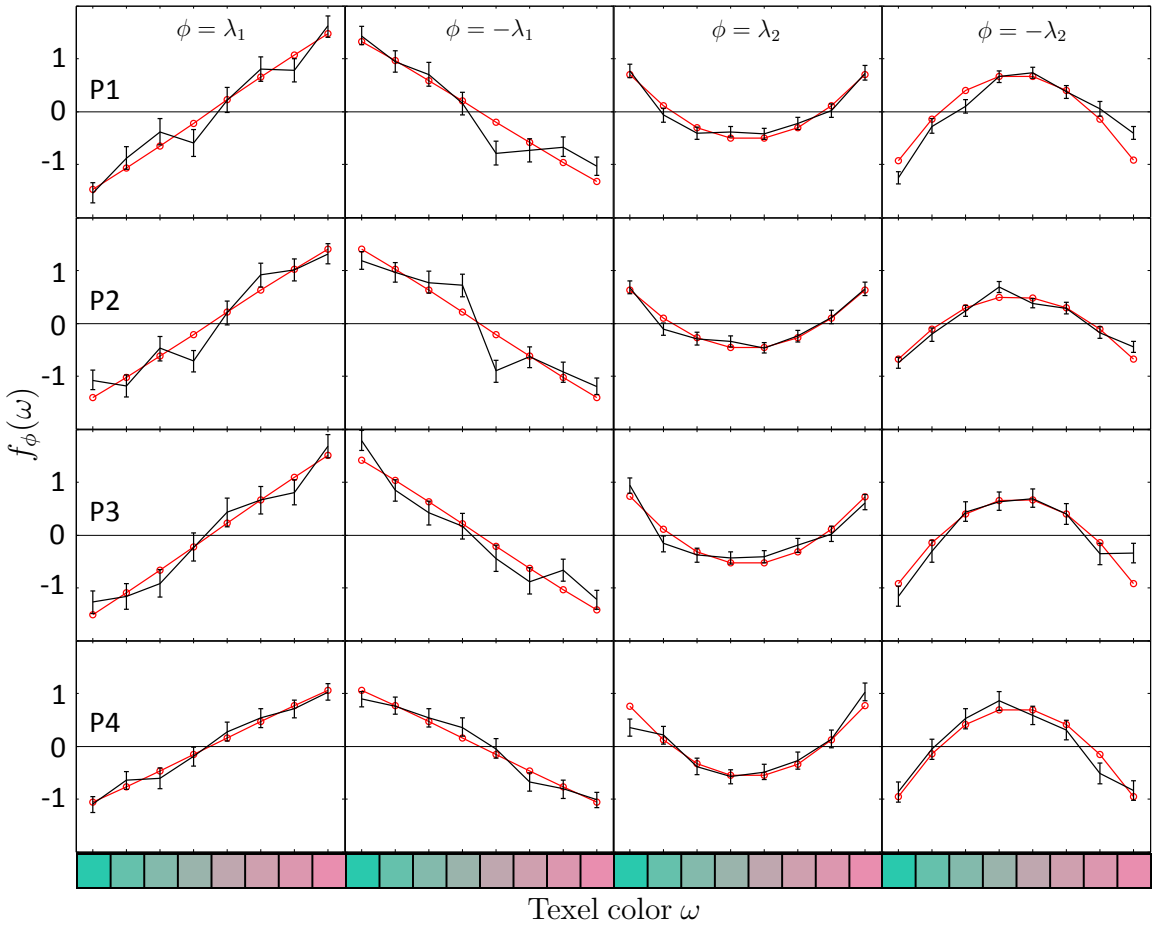

Figure 3.3: *Expansions achieved by all participants in all conditions.* For seeds $\phi = \lambda_1$, $-\lambda_1$, λ_2 and $-\lambda_2$ (in panels from left to right) are plotted the corresponding expansions f_ϕ achieved by participants P1, P2, P3 and P4 (in panels from top to bottom). The red curve in each panel shows the seed function ϕ scaled to have norm equal to the norm of the expansion. Note that in nearly all cases the expansion closely matches the seed.

The main thing to note is that for each participant in each seed condition, the expansion f_ϕ matches very closely the seed perturbation ϕ; what do we take away from this? Several points are relevant:

1. We plot the expansion f_ϕ achieved by a given participant in the condition with seed perturbation ϕ as having a mean of 0; however, this does not mean that the attention filter used by the participant to perform the task actually has a mean of 0. Our experimental design enables us to determine this filter only up to an unmeasurable additive constant. Plausibly, in fact, the activation of any given mechanism is signaled by positive deviations from a resting state of 0. If so, then the actual attention filter used by the participant in any condition must assign only nonnegative values.

2. Our task is designed to encourage the participant to use top-down attention in the condition with seed ϕ to fashion an attention filter that gives high activation to the target but not to the background. This implies that a mechanism with sensitivity function f will only be useful for detecting the target in the condition with seed perturbation ϕ if $f \bullet \phi > 0$. (Otherwise, this mechanism will be more strongly activated by the background than by the target.) Note the implication that if a mechanism is useful for detecting the target in the condition with seed perturbation ϕ, then it will not be useful for detecting the target in the complementary condition with seed perturbation $-\phi$.

Because participants achieve expansions f_ϕ that closely match the form of ϕ for both $\phi = \lambda_1$ and also for $\phi = -\lambda_1$, it follows that they must possess mechanisms whose sensitivity functions are correlated positively with λ_1 as well as mechanisms whose sensitivity functions are correlated negatively with λ_1. Similarly, because participants achieve expansions f_ϕ that closely match the form of ϕ for both $\phi = \lambda_2$ and also for $\phi = -\lambda_2$, they must possess mechanisms whose sensitivity functions are correlated positively with λ_1 as well as other

mechanisms whose sensitivity functions are correlated negatively with λ_2. We conclude that human vision must possess at least four distinct mechanisms differentially sensitive to the colors in Ω.

On the other hand, for $\phi = \lambda_1$, the expansion f_ϕ is nearly perfectly complementary in form to $f_{-\phi}$. This suggests that whatever mechanisms the participant is using to perform the task for $\phi = \lambda_1$, he/she has mechanisms with complementary sensitivity functions that can be used to perform the task for $\phi = -\lambda_1$.

3.3.1 The model

The observations above suggest that it may be possible to describe the results in Fig. 3.3 with the following model:

1. There exist two functions $h_1 : \Omega \to \mathbb{R}$ and $h_2 : \Omega \to \mathbb{R}$, each summing to 0 and satisfying $\|h_1\| = \|h_2\| = 1$, such that for a given participant j, the mechanisms useful for discriminating Ω-scrambles have sensitivity functions

$$F_{j,1}(\omega) = C_{j,1} + A_{j,1}h_1(\omega) \quad \text{and} \quad F_{j,2}(\omega) = C_{j,2} - A_{j,2}h_1(\omega) \tag{3.8}$$

$$F_{j,3}(\omega) = C_{j,3} + A_{j,3}h_2(\omega) \quad \text{and} \quad F_{j,4}(\omega) = C_{j,4} - A_{j,4}h_2(\omega) \tag{3.9}$$

2. In the condition with a given seed ϕ, the Ω-filter $F_{\phi,j}$ achieved by participant j is given by

$$F_{\phi,j} = \sum_{k=1}^{4} w_{\phi,j,k} F_{j,k} \tag{3.10}$$

for which the weights $w_{\phi,j,1}, w_{\phi,j,2}, w_{\phi,j,3}, w_{\phi,j,4}$ maximize the difference in activation produced by Ω-scrambles with histograms $U + \phi$ vs $U - \phi$, under the constraints that

(a) the weights are all nonnegative;

(b) the weights sum to 1.

3. Probability correct in the condition with seed ϕ on a trial in which the perturbation tested was ρ is $\Psi_j(\mathrm{Sal}_{\phi,j})$, where

$$\mathrm{Sal}_{\phi,j}(\rho) = F_{\phi,j} \bullet \rho \quad \text{and} \quad \Psi_j(x) = 0.125 + 0.855\left(1 - \exp\left[-x^{\beta_j}\right]\right), \quad x \in \mathbb{R}^+. \tag{3.11}$$

Because each of the functions h_1, h_2 is constrained to sum to 0 and to have norm equal to 1, these functions collectively contribute $2 \times 6 = 12$ degrees of freedom. The model is invariant with respect to the additive scalars $C_{j,k}$, so they add no degrees of freedom to the model. Each of the parameters β_j and $A_{j,k}$, $j = 1, 2, 3, 4$, $k = 1, 2, 3, 4$ adds a degree of freedom. Thus the total number of degrees of freedom is 32.

3.3.1.1 Results of fitting the model

Fig. 3.4 shows the four estimated mechanism sensitivity functions $F_{j,k}(\omega)$, $k = 1, 2, 3, 4$, for participants Pj, $j = 1, 2, 3, 4$, from left to right. Only the 0-mean functions $A_{j,k} f_k(\omega)$ have actually been estimated; for plotting purposes we set $C_{j,k} = -\min\{f_{j,k}\}$ in each case to make $\min\{F_{j,k}\} = 0$. Note that (by construction) sensitivity functions $F_{j,1}$ and $F_{j,2}$ are linearly dependent as are sensitivity functions $F_{j,3}$ and $F_{j,4}$.

Fig. 3.5 plots (in red) the expansions predicted by the model for all four participants. The black curves reproduce the black curves from Fig. 3.3; these are the expansions estimated from the data for each individual subject in each seed condition separately. The number of degrees of freedom used to produce the black (red) curves in Fig. 3.5 is $7 \times 16 = 112$ (32). However, the red curves account for more than 98% of the variance in the trial-by-trial

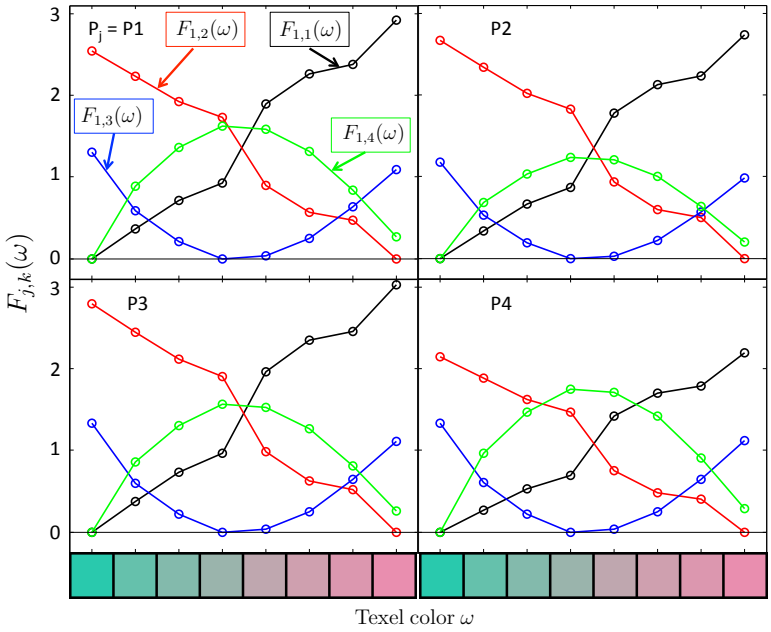

Figure 3.4: *Estimated mechanism sensitivity functions.* Fitting the model (using a maximum likelihood criterion) to the data for all four participants $j = 1, 2, 3, 4$ yields the four estimated sensitivity functions $F_{j,k}$, $k = 1, 2, 3, 4$. In each case, $F_{j,k}$ has been shifted vertically to make its minimum value equal to 0. Results are shown for participants P1, P2, P3 and P4 in separate rows of panels. Sensitivity functions $F_{j,1}$ and $F_{j,2}$ are constrained by the model to be linearly dependent; the same is true of $F_{j,3}$ and $F_{j,4}$.

saliences (across all 41,600 trials) predicted using the expansions (the black curves) estimated separately for all participants in all seed conditions.

3.4 Discussion for Chapter 3

We used the "seed expansion" method to investigate what mechanisms are sensitive to colored Ω-scrambles, a carefully crafted texture that isolates visual mechanisms. To confine the posible mechanisms to be sensitive only to changes in color, we created an indivualized color palette of motion-equiluminant lights using the method described in Chapter 1. Our results, plotted in Fig. 3.3 indicate that human vision has at least 4 distinct mechanisms differentially sensitive to Ω-scrambles. One mechanism confers sensitivity that roughly increases linearly across the gamut of Ω from red to green. A second mechanism, complementary to the first,

confers sensitivity that increases roughly linearly across the gamut of Ω from green to red. A third mechanism confers sensitivity that increases roughly evenly with increasing red/green saturation, and the fourth mechanism, complementary to the third, confers sensitivity that increases roughly evenly with decreasing red/green saturation. One of our goals was to provide evidence concerning the existence of half-axis mechanisms. Our results provide no clear evidence for the existance of such half-axis. Specifically, none of the mechanisms we describe takes the form of a half-axis mechanism. We would like to point out that it is not clear that the current experiment provides a strong tests fo this conjecture, and the existence of half-axis mechanims is therefore not precluded based on the current experiments.

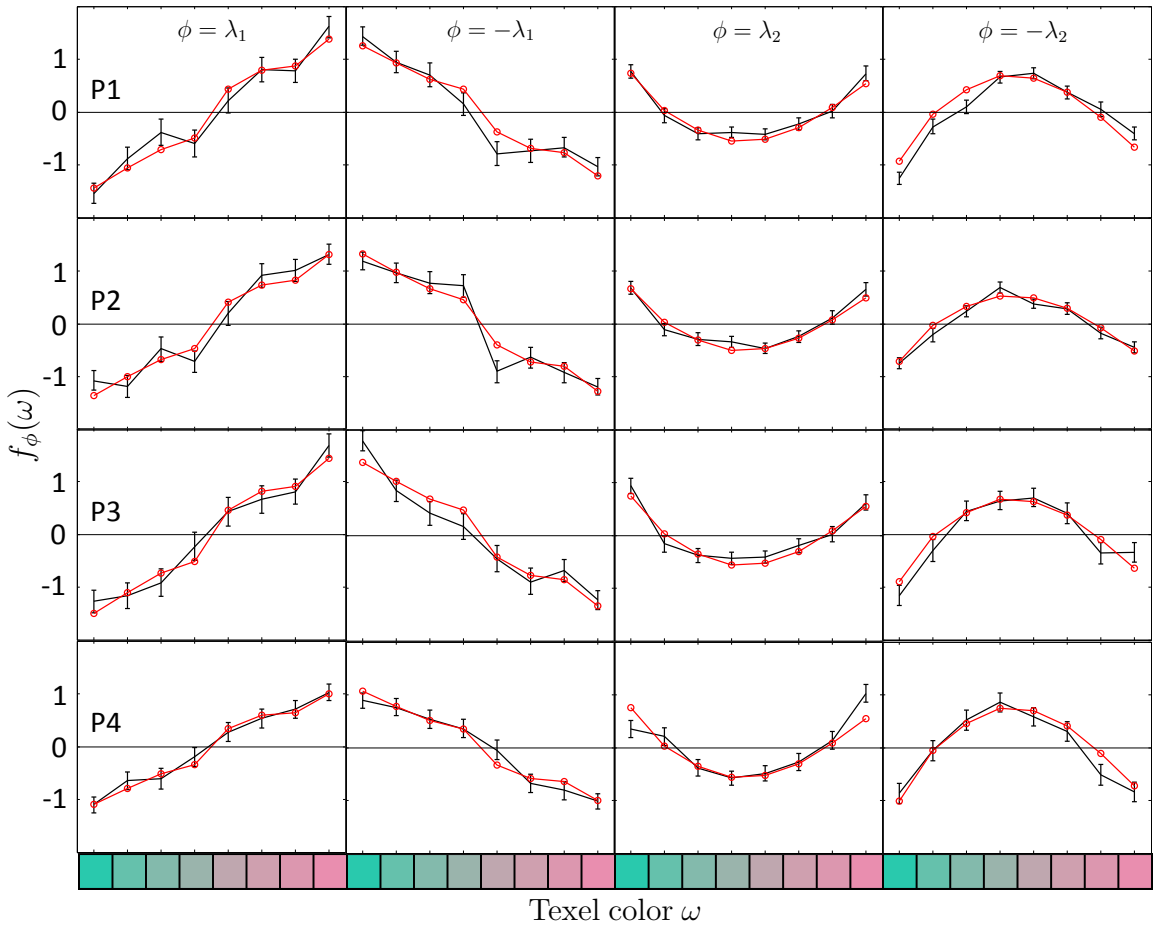

Figure 3.5: *Expansions predicted by the model.* Expansions estimated from the model (plotted in red) and expansions estimated from the data from individual seed conditions (plotted in black) for each of the three participants. Error bars are 95% Bayesian credible intervals. Note that the model expansions (red curves–based on 32 degrees of freedom) account for more than 98% of the variance in the trial-by-trial saliences (across all 41,600 trials) estimated using the expansions (the black curves) derived separately for all participants in all seed conditions.

Bibliography

E. H. Adelson and J. R. Bergen. Spatio-temporal enegry models for the perception of apparent motion. *Journal of the Optical Society of America A*, 2:284–299, 1985.

S. Anstis and P. Cavanagh. A minimum motion technique for judging equiluminance. In J. D. Mollon and E. T. Sharpe, editors, *Colour Vision*, pages 155–166. Academic Press, 1983.

R. M. Boynton and P. K. Kaiser. Vision: The additivity law made to work for heterochromatic photometry with bipartite fields. *Science*, 161(3839):366–368, 1968.

D. H. Brainard. The psychophysics toolbox. *Spatial Vision*, 10:433–436, 1997.

C. Chubb, I. Scofield, C-C. Chiao, and G. Sperling. A method for analyzing the dimensions of preattentive visual sensitivity. *Journal of Mathematical Psychology*, 56:427–443, 2012.

R. T. Eskew. Higher order color mechanisms: A critical review. *Vision Research*, 49:2686–2704, 2009.

R. T. Eskew, J. R. Newton, and F. Giulianini. Chromatic detection and discrimination analyzed by a bayesian classifier. *Vision Research*, 41:893–909, 2001.

K. R. Gegenfurtner and D. C. Kiper. Contrast detection in luminance and chromatic noise. *Journal of the Optical Society of America A*, 9:1880–1888, 1992.

P. D. Gowdy, C. F. Stromeyer, and R. E. Kronauer. Detection of flickering edges: absence of a red-green edge detector. *Vision Research*, 39(25):4186–4191, 1999a.

P. D. Gowdy, C. F. Stromeyer, and R. E. Kronauer. Facilitation between the luminance and red-green detection mechanisms: enhancing contrast differences across edges. *Vision Research*, 39(24):4098–4112, 1999b.

N. Graham. *Visual Pattern Analyzers*. Oxford University Press, New York, 1989.

H. E. Ives. Studies of the photometry of lights of different colours. i. spectral luminosity curves obtained by the equality of brightness photometer and the flicker photometer under similar conditions. *Philosophical Magazine*, 24:149–188, 1912a.

H. E. Ives. Studies of the photometry of lights of different colours. ii. spectral luminosity curves by the method of critical frequency. *Philosophical Magazine*, 24:352–370, 1912b.

H. E. Ives. Studies of the photometry of lights of different colours. iii. distortions in spectral luminosity curves produced by variations in the character of the comparison standard and of the surroundings of the photometric field. *Philosophical Magazine*, 24:744–751, 1912c.

H. E. Ives. Studies of the photometry of light of different colours. iv. the addition of luminosities of different color. *Philosophical Magazine*, 24:845–853, 1912d.

H. E. Ives. Studies of the photometry of light of different colours. v. the spectral luminosity of the average eye. *Philosophical Magazine*, 24:853–863, 1912e.

H. E. Ives. The theory of the flicker photometer. *Philosophical Magazine*, 28:708–728, 1914.

H. E. Ives. A polarization flicker photometer and some data of theoretical bearing obtained with it. *Philosophical Magazine*, 33:360–380, 1917.

J. Krauskopf and K. R. Gegenfurtner. Color discrimination and adaptation. *Vision Research*, 32:2165–2175, 1992.

J. Krauskopf and Q. Zaidi. Induced desensitization. *Vision Research*, 26(5):759–762, 1986.

J. Krauskopf, D. R. Williams, and D. W. Heeley. Cardinal directions of color space. *Vision Research*, 26:23–32, 1982.

J. Krauskopf, D. R. Williams, D. W. Heeley, and A. M. Brown. Higher order color mechanisms. *Vision Research*, 36:1235–1245, 1986.

J. Krauskopf, H. Wu, and B. Farell. Coherence, cardinal directions and higher-order mechanisms. *Vision Research*, 36:1235–1245, 1996a.

J. Krauskopf, Q. Zaidi, and M. B. Mandler. Mechanisms of simultaneous color induction. *Journal of the Optical Society of America A*, 3:1752–1757, 1996b.

Z.-L. Lu and G. Sperling. The functional architecture of human visual motion perception. *Vision Research*, 35:2697–2722, 1995.

Z.-L. Lu and G. Sperling. Sensitive calibration procedures based on the amplification principle in motion perception. *Vision Research*, 41:2355–2374, 2001.

K. T. Mullen and J. J. Kulikowski. Wavelength discrimination at detection threshold. *Journal of the Optical Society of America A*, 7:733–742, 1990.

M. J. Sankeralli and K. T. Mullen. Bipolar or rectified chromatic detection mechanisms? *Visual Neuroscience*, 18(1):127–135, 2001.

L. T. Sharpe, A. Stockman, W. Jagla, and H. Jagle. Photopic and mesopic photometry: yesterday, today and tomorrow. *Journal of Vision*, 5(11):1–12, 2005.

A. E. Silva and C. Chubb. The 3-dimensional, 4-channel model of human visual sensitivity to grayscale scrambles. *Vision Research*, 101:94–107, 2014.

J. P. H. van Santen and G. Sperling. Temporal covariance model of human motion perception. *Journal of the Optical Society of America A*, 1:451–473, 1984.

J. P. H. van Santen and G. Sperling. Elaborated reichardt detectors. *Journal of the Optical Society of America A*, 2:300–321, 1985.

G. Wagner and R. M. Boynton. Comparison of four methods of heterochromatic photometry. *Journal of the Optical Society of America*, 62(12):1508–1515, 1972.

A. B. Watson and A. J. Ahumada. A model of human visual-motion sensing. *Journal of the Optical Society of America A*, 2:322–342, 1985.

A. B. Watson and J. G. Robson. Discrimination at threshold: Labelled detectors in human vision. *Vision Research*, 21:1115–1122, 1981.

M. A. Webster and J. D. Mollon. Changes in colour appearance following post-receptoral adaptation. *Nature*, 349:235–238, 1991.

M. A. Webster and J. D. Mollon. The influence of contrast and adaptation on color appearance. *Vision Research*, 34:1993–2020, 1994.