

University of Nevada, Reno

Counting Colors: Judgments of Summary Statistics in Color Arrays

A thesis submitted in partial fulfillment of the
Requirements for the degree of
Master of Arts in Psychology

By

Jacquelyn Webster

Dr. Michael A. Webster/Thesis Advisor

December, 2013



University of Nevada, Reno
Statewide • Worldwide

THE GRADUATE SCHOOL

We recommend that the thesis
prepared under our supervision by

JACQUELYN WEBSTER

entitled

Counting Colors: Judgments Of Summary Statistics In Color Arrays

be accepted in partial fulfillment of the
requirements for the degree of

MASTER OF ARTS

Michael A. Webster, Advisor

Gideon P. Caplovitz, Committee Member

Louis C. Forline, Graduate School Representative

Marsha H. Read, Ph. D., Dean, Graduate School

December, 2013

Abstract

The visual system can reliably extract the average value of stimulus distributions that vary along many dimensions, from orientation to facial expressions. We examined how sensitive observers were to the average chromaticity of color distributions. This average has special significance for color coding (e.g. to estimate the illuminant) but how it might depend on the level of representation (e.g. perceptual vs. cone-opponent) or nonlinearities (e.g. categorical coding) is unknown. Stimuli were arrays of spots that alternated between two component colors with differing hues that bracketed a given color boundary. Observers set the array so that the composite color equaled a specified reference hue. These settings could be made reliably though variance increased with increasing angular difference between the components. Matches for an average red, blue-red, or yellow-green were consistent with the arithmetic mean chromaticity of the pair, while blue-green settings instead deviated toward blue (indicating greater perceptual weight for the green component). The settings show little evidence for categorical coding of the colors, and moreover cannot be predicted from the scaled red-green and blue-yellow appearances of the individual components, suggesting that there is little influence of linguistic or phenomenal color categories on the mixture judgments.

Keywords: color vision, psychophysics, ensemble perception

Acknowledgments

Supported by EY-10834

Table of Contents

Abstract.....	i
Acknowledgments.....	ii
Table of Contents.....	iii
Figure Captions.....	iv
Introduction.....	1
Methods.....	5
Participants.....	5
Display and Stimuli.....	6
Procedure.....	7
Results.....	9
Yoke Hue Pairs.....	9
Hue Pairs With A Fixed Component.....	15
Discussion.....	19
References.....	24
Appendices.....	27

Figure Captions

- Figure 1. Stimulus Array: An example of the stimulus array with the elements alternating between the two colors. In one task (yoked hue pairs) the hue angle between the two components remained constant and the pair were varied together to make the setting. In the second task (fixed single component) the color corresponding to the corner elements remained fixed while observers varied the hue of the second component.
- Figure 2. Yoked Hue Pairs: Examples of the settings made by two individual observers for each of the 4 reference hues (BG blue-green; RB: blue-red; RB; red: R; YG: yellow-green). Filled symbols plot the settings when both components had the same hue (i.e. a hue-angle difference of 0) and presented at four levels of contrast. Unfilled symbols plot the settings when the pair differed in hue angle by 10, 20, 40, or 80 deg, with larger differences plotted closer to the origin consistent with the actual chromaticity of the mixture.
- Figure 3a. Yoked Hue Pairs: Mean blue-green settings for 13 observers as a function of the contrast of a uniform-hue array (leftmost points) or the hue-angle difference of pairs with a fixed contrast of 80 (rightmost points). Error bars = ± 1 SEM.
- Figure 3b. Yoked Hue Pairs: Average standard deviation of the individual's settings for the same stimuli, based on 4 repeated settings per condition. Error bars = ± 1 SEM. Dashed lines – predicted difference in the standard deviation of the settings for the uniform vs. 2-component arrays if the errors in judging the two components are independent and additive.
- Figure 4. Yoked Hue Pairs: Average settings for the 4 hues, based on 4 or more observers tested per condition. As in Figure 2, filled circles plot the settings for uniform arrays presented at four levels of contrast. Unfilled circles plot the settings for 2-component arrays where contrast was fixed at a value of 80. Gray squares show the settings predicted if observers are basing their responses on the mean of the log cone-opponent responses (top) or on the average of the perceived red-green and blue-yellow proportions in each component based on hue scaling (bottom). Italic characters in the bottom figure show the loci of the unique and binary hues given by the hue-scaling function.
- Figure 5. Fixed Single Component: Blue-green settings for 4 individual observers when one of the two color components remained fixed. Filled symbols plot the original measured settings. Unfilled symbols plot the mirror settings by exchanging the abscissa and ordinate. Gray lines plot the settings predicted by averaging the linear or log cone-opponent responses.

Figure 6. Fixed Single Component: Predicted settings for different levels of categorical bias, assuming the unbiased settings are based on the arithmetic mean of the hue angles. Settings range from a slope of -1 for no categorical bias to constant settings at the blue-green boundary for a complete bias.

Figure 7. Fixed Single Component: Comparisons between the original fixed-hue settings and the corresponding setting when the variable hue was assumed fixed. Positive diagonal corresponds to equal symmetric matches in the two cases. Symbols plot the settings for individual observers. Each panel shows the settings for the labeled average hue.

Counting Colors: Judgments of Summary Statistics in Color Arrays

Introduction

Recently a wide array of studies have explored the perception of “ensemble statistics” of scenes, such as the ability to judge the average value of stimuli composed of distributed features (D. Whitney, Haberman, & Sweeny, In Press). Early reports showed that observers can accurately estimate the average direction of motion of a field of randomly moving dots (Watamaniuk, Sekuler, & Williams, 1989; Williams & Sekuler, 1984), the average orientation of a set of lines (Dakin, Bex, Cass, & Watt, 2009; Robitaille & Harris, 2011), or the average size of a set of circles (Ariely, 2001). Moreover, these judgments can be made without attention (Alvarez & Oliva, 2008) or when elements in the display are rendered invisible, for example by crowding (Parkes, Lund, Angelucci, Solomon, & Morgan, 2001). The average of a stimulus set can also be represented more accurately than the individual members, and observers may be sensitive to changes in the average even when they are unaware of a change in the individual elements (Haberman & Whitney). This has led to the suggestion that the visual system might directly encode summary statistics such as the average feature properties in order to allow efficient representations of scenes. Similar summary representations have been found for higher-level properties of images, including the perceived gender, identity, or expression of faces (Haberman, Harp, & Whitney, 2009; Haberman & Whitney, 2007, 2009), or the average direction of point-light walkers simulating biological motion (D. Whitney, Haroz, & Whitney, 2013). They have also been shown for other modalities such the perception of

sound ensembles (J. H. McDermott, Schemitsch, & Simoncelli, 2013; Piazza, Sweeny, Wessel, Silver, & Whitney, 2013). This suggests that ensemble coding is a common and consistent property of sensory representations.

In this study we explored ensemble judgments for color. Despite the recent spate of interest in summary statistics in perception, there has been little corresponding work within the context of color perception. Yet for color the mean may play a special role. For example, the average color in a scene is a clue to the illuminant and thus an important potential cue for color constancy, allowing the visual system to discount spectral variations in order to assign stable colors to objects (Foster, 2011). Moreover, mechanisms sensitive to the mean color occur very early in the visual pathway. The receptors themselves adapt to the average color sampled over time by successively scanning the scene, a process that is again fundamental to color constancy (Brainard & Wandell, 1992; D'Zmura & Lennie, 1986). Individuals are also very sensitive to a common vs. non-uniform shift in the chromaticities of a set of colors, and thus can readily detect sudden changes in the mean (Foster & Nascimento, 1994).

How colors might be perceptually averaged can potentially reveal the stages and representations on which this averaging depends. Color coding is known to undergo a series of transformations at different levels of the visual system (Conway et al., 2010). Within the retina and LGN, the initial signals from the long, medium, and short-wavelength sensitive cones (L, M, and S) are re-organized in terms of three cardinal cone-opponent axes: two that carry chromatic contrast based on comparing either the L and M cones (LvsM) or the

S signals opposed by the combined L and M signals (SvsLM), and a third non-opponent dimension conveying lightness or luminance (Derrington, Krauskopf, & Lennie, 1984). At cortical levels these are elaborated into “higher-order” mechanisms that include color channels tuned to a variety of different hue and lightness directions, and these may not retain the separable contrasts along the cardinal dimensions (Krauskopf, Williams, Mandler, & Brown, 1986; Lennie, Krauskopf, & Sclar, 1990). For example, cortical responses and visual performance on some tasks may be weaker for the blue-yellow dimension of color space than for a purplish-greenish direction that has the same LvsM and SvsLM component contrasts, but paired in opposite phase (Goddard, Mannion, McDonald, Solomon, & Clifford, 2010; Juricevic, Land, Wilkins, & Webster, 2010; K. C. McDermott, Malkoc, Mulligan, & Webster, 2010). The number of further potential transformations remains unclear. However, it is well known that the red vs. green and blue vs. yellow perceptual opponencies that are thought to define the principal dimensions of our phenomenal experience of color are not aligned with the cardinal axes (De Valois, De Valois, Switkes, & Mahon, 1997; Krauskopf, Williams, & Heeley, 1982; Webster, Miyahara, Malkoc, & Raker, 2000; Wuerger, Atkinson, & Cropper, 2005). For example, as suggested above, stimulus variations along a perceptually pure blue-yellow axis do not isolate either of the cardinal chromatic axes and are instead roughly midway between them (i.e. these colors lie along a diagonal in the plane defined by the cardinal chromatic axes). Moreover, opposing unique hues such as red and green are not complementary (on opposite sides of gray) in terms of the

cardinal axes. Such differences raise the question of which if any “space” an observer might operate within to estimate the relationships or average of a set of colors.

Finally, the representation of color also includes a number of potential nonlinearities that could influence averaging judgments. For example, cone signals and opponent-cone contrasts may reflect compressive nonlinearities in the neural responses which could bias the linear mean of a set of contrasts (e.g. (Horwitz & Hass, 2012; Ruderman, Cronin, & Chiao, 1998)), though in practice responses to moderate contrasts appear roughly linear (Webster & Wilson, 2000). At higher levels, color coding may also be impacted by nonlinearities such as categorical coding, in which the stimulus is represented discretely (e.g. as simply green or blue) rather than as a continuously varying hue. Categorical effects have been reported for color in a number of studies (e.g. (Franklin et al., 2008; Gilbert, Regier, Kay, & Ivry, 2006; Roberson, Pak, & Hanley, 2008; Winawer et al., 2007)), though they remain controversial and it is unclear under which conditions or tasks they might be manifest (Brown, Lindsey, & Guckes, 2011; Webster & Kay, 2012; Witzel & Gegenfurtner, 2011). A tendency to categorize when averaging colors could again lead to different, nonlinear weightings of the component hues in the stimulus, which should again bias estimates of the average in predictable ways.

Summary judgments of color distributions thus provide a potentially rich context for exploring not only ensemble perception but also many of the properties of color perception itself. In the current work, we examine these

summary estimates using pairs of perceptually similar colors (e.g. blue and green), where the task of the observer is to judge the relative balance of the pair. A companion paper explores our ability to judge the metrical relationships between more disparate color categories. Here we focus on how well observers can actually judge the mean of colors, whether these means reflect nonlinearities in color coding, and which level of representation (e.g. cone-opponent or perceptual) might underlie the judgments.

Methods

Participants

Observers included one of the authors (JW) and 16 additional participants who were University of Nevada, Reno students and were unaware of the theoretical aims of the study. Different observers participated in two different subsets of experiments consisting of a yoked hue pair condition and a fixed single component condition. Both included a uniform hue condition where the array consisted of a single hue to be used as a baseline comparison. 15 participants were included in the yoked hue pair condition, 13 of which were tested on blue-green colors in order to fully evaluate the settings for this color pair, since it has been a focus of categorical color coding. Smaller numbers were tested on the remaining color pairs; 3 for both green-yellow and red-blue, and 5 for red. 9 participants were included in the fixed single component condition, 7 of which who had participated in the previous yoked hue condition. Again, a larger number of participants were tested on blue-green colors (7 out

of the 9) and 4 of the 9 participants in this condition were tested on each of the remaining color pairs. All participants had normal color vision as assessed by the Cambridge Colour Test and normal or corrected-to-normal acuity.

Participation was with informed consent following protocols approved by the university's Institutional Review Board.

Display and Stimuli

Stimuli were presented on a Sony 500PS CRT monitor controlled by a Cambridge Research Systems VSG 2/5 graphics card, which allows colors to be displayed with high resolution. The monitor was calibrated and gamma-corrected using a Photo Research PR650 spectroradiometer.

For all experiments the stimulus consisted of an 11 by 11 square array of 0.75deg uniform circles with a center-to-center spacing of 0.02deg (Figure 1). The array was shown on a uniform 15 by 20 deg gray background with the chromaticity of (CIE 1931 x,y: 0.31, 0.316) and the same luminance as the array (30cd/m²). Each circle included a narrow black outline to help demarcate it from the background. The array could be composed of a single color common to all elements, or alternated between two different colors. The colors were specified within a scaled version of cardinal axis space, in which the two axes correspond to the LvsM and SvsLM cardinal chromatic directions and chromaticities are given by their angle and distance from the gray neutral point. Signals along the LvsM and SvsLM axes were scaled based on previous studies to roughly equate contrast sensitivity with units corresponding very roughly to multiples of threshold. The relationship between coordinates in our

scaled space and the r_{mb} , b_{mb} coordinates of the MacLeod-Boynton chromaticity diagram (MacLeod & Boynton, 1979) is given by:

$$\text{LvsM contrast} = (r_{mb} - 0.6568) * 2754 \quad (1)$$

$$\text{SvsLM contrast} = (b_{mb} - 0.01825) * 4099 \quad (2)$$

A common scaling was used for all observers, with luminance defined photometrically. Contrast matches across the different axes (not reported) showed that this nominal scaling roughly equated the perceived suprathreshold contrasts for our observers. In uniform hue conditions contrast was varied at values of 20, 40, 60, or 80 while the hue angle of the components remained the same (in other words, there was no hue difference between the components). For the experiments, the contrast was fixed at a value of 80, and the component colors were instead varied only in their hue angle. For the cardinal axes these angles corresponded to 0 and 180 deg (+LvsM and -LvsM) and 90 and 270 deg (+S and -S).

Procedure

Observers freely viewed the display binocularly from a distance of 114cm in an otherwise dark room. In some trials (yoked hue pairs) the angular difference between the pair of colors remained fixed while observers adjusted the absolute angles of the pair. In a second set of conditions (fixed single component) one of the component colors was fixed and the observer instead varied the angle of the remaining color. In both cases the task of the observer was to adjust the pair until the average matched a specified criterion. This

specified criterion color was never shown to the observer in order to rule out the possibility of the results being reflective of a simple color matching task and to determine if observers could reliably extract the average value of a set of different colored stimuli. Specified criteria included balanced binary hues of blue-green (i.e. an equal perceived mixture of blue and green), red-blue, or green-yellow, or a unique red (i.e. a red that appeared neither bluish nor yellowish). These four hues were selected because for single colors they lie roughly along the 4 poles of the cone-opponent space, and thus allowed us to sample how colors were averaged in terms of the signals along the cardinal axes. Additionally four levels of hue angle distance were tested in both conditions, for each of the four specified criteria, including 10, 20, 40, and 80 deg separations. In the yoked hue pair condition, the degree of angle separation remained fixed throughout trials whereas in the fixed single component condition, the test stimuli began at these different degrees of angle separation but was varied by the observer throughout each trial.

The instructions to the observer were, for example in the blue-green trials, to 'Adjust the pair of colors until the average appears blue-green' or so that there is 'an equal amount of blue and green present' if they needed further instruction. Before making adjustments at the beginning of each trial, the hue angles of the test stimuli were presented at random hue angles around the cardinal axes of each of the four specified criteria. Hue angles of the test stimuli were then varied with a staircase procedure based on the response to the pair (e.g. "too blue" or "too green" for the blue-green judgment) on each

trial, with the chosen angle based on the mean of the last 4 reversals and the array was presented until a judgment was made. Typically 4 settings were made for each hue combination, with the angle separation between the components counterbalanced, so that each initial combination of bracketing pairs and hue criterion was presented four times for each observer's adjustments. Finally, standard deviation was obtained for those four presentations and then averaged across observers.

Results

Yoked Hue Pairs

In the first set of experiments we explored the average hue in color pairs that differed by a fixed hue angle and thus were covaried to set the average. Figure 2 illustrates the resulting settings for two representative observers, by plotting the coordinates of the settings in the LvsM and SvsLM plane. The filled symbols show the settings when the array was composed of a single color, with the contrast fixed at a value between 20 and 80. These replicate the basic features found previously for the loci of unique and binary hues in the cone-opponent plane (Webster et al., 2000; Wuerger et al., 2005). Unique red is close to the +L pole of the LvsM axis, while the remaining three poles instead correspond to balanced mixtures of blue-green (-L), red-blue (+S), or yellow-green (-S). The two figures also illustrate the large individual differences in hue loci among color-normal observers (Kuehni, 2004; Webster et al., 2000).

The unfilled symbols plot the settings when the array was composed of two different hues, with the difference varying between 10 and 80 degrees, with the contrast of the two components fixed at 80. To aid visualization, the contrast of the average is arbitrarily given by the actual mean of the two chromaticities in the pair, which obviously plots within the circle on which the two components lie, and closer to the origin the greater their separation. For example, for a hue difference of 80 the mixture has a contrast of 61. (Importantly, however, this is not intended to imply that for the observers the average for the pair necessarily appeared to have a lower average contrast than either component.)

For both observers there are conspicuous biases in the blue-green settings as the average difference between the two component hues increases, and this was a consistent trend across all observers. For example, Figure 3a shows the mean settings for 12 individuals. As the angular difference within the pairs increases, the average rotates toward progressively shorter angles. This corresponds to assigning less weight to blue (or more weight to green) Note that this effect cannot reflect changes in the contrast of the physical mean, since the hue angles chosen with the uniform arrays remain roughly independent of contrast. Instead, there is a strong and consistent nonlinearity in blue-green average judgments, at least when represented within the linear cone-opponent space.

The mean shifts with increasing separation were accompanied by more variable judgments of the average. This is shown in Figure 3b, which plots the

average standard deviation of the 12 observers' settings. Variance was significantly greater even when the two components differed only by 20 deg ($t = 2.91$, $df = 22$, $p = .008$). If sensitivity to the average were limited by the uncertainty in judging the hue of the individual components, then the variance should be twice as large for the pairs (Figure 3b, dashed line). However, by a separation of 80 the estimates exceed this value ($t = 3.50$, $df = 11$, $p = .005$). Thus the ability to estimate the average of the two colors shows both systematic biases and increasing error as the difference between the colors increases.

While individual differences were large, we did not find similar consistent biases for the other 3 reference hues examined. This can be seen in Figure 4, which plots the average of the settings for all observers for each of the 4 measured hue loci. Average settings for red remained aligned with the uniform color judgments along the +L axis. The averages for blue-red and yellow-green are instead shifted slightly though again there were large individual differences and thus these overall shifts are not significant.

We considered four possible bases for the observed matches. The first was optical, and assumed that the matches are linear with cone excitations but appear biased in the blue-green region because of chromatic aberration. Wavelength-dependent blurring produces large losses in chromatic contrast in the retinal image, and is a major factor contributing to the fall in chromatic contrast sensitivity with increasing spatial frequency (Poirson & Wandell, 1993). These aberrations may also be responsible for large size-dependent changes

in the relative salience of different chromatic signals in arrays similar to those we used (Regan & Mollon, 1997). Obviously, simply blurring the two colors together would not be expected to alter the average. Yet chromatic aberration could impact the balance if it blurred one component more than the other, and the predicted differences in blur are in fact largest for the blue and green pair (Howarth, 1984). Thus we cannot exclude some contribution of this factor.

However, it is unlikely to be the only cause of the biases. First, the fundamental spatial frequency of the array was in fact low (0.5 c/deg). Second, if the matches are based on simply the average chromaticity, then the required contrast differences between the blue and green are far too large (e.g. assuming the blue green balance is at 180 deg, the match at 160 deg when the blue and green are 80 deg apart requires that the blue contrast is only 40% of the green contrast: $\sin(200^\circ) + 0.4\sin(120^\circ) = \sin(180^\circ)$). Instead, direct contrast matches between stimuli at different hue angles suggested that perceived contrast was similar along the different axes. Finally, it cannot be due to blur-induced changes in the perceived hue angle of the stimuli, since the perceived hue of the spots remained stable when we doubled or halved the viewing distance.

The foregoing suggests that at least part of the blue-green bias reflects non-linearities in the cone-opponent responses. For example, suppose that observers' judgments are based on a compressive nonlinearity such as the log of the cone responses (Figure 4a). This would predict that S cone increments (hue angles >0 deg and < 180 deg and thus above the LvsM axis) would have

weaker effective contrast than S cone decrements (angles >180 deg and <360 deg). Predictions based on this assumption (gray squares in Figure 4a) provide a rough approximation to the shifts toward (weaker) blues in the blue-green settings. It also predicts the absence of biases along the S axis itself, where the average instead depends on balancing the L and M responses. This is because along this axis the L and M signals are opposed, and thus the modeled nonlinearities within each largely cancel. This asymmetry is seen in how light adaptation alters the effective color contrast along the LvsM and SvsLM axes (Webster & Mollon, 1995; Zaidi, 1998). However, this account fails for the red settings, since like the blue-green axis these again depend primarily on balancing the S cone signals and thus should show the same biases as blue-green. Moreover, nonlinearities in the cone-opponent contrasts are more generally an unlikely explanation because as we noted above, under a constant state of adaptation the cone-opponent contrasts appear fairly linear (Derrington et al., 1984; Switkes, 2008; Webster & Wilson, 2000).

In the third case, we asked whether the average hues could be predicted from the average *appearance* of the component colors. That is, observers might first judge the proportion of blue and green in the individual elements, and then adjust these so that these perceived proportions are balanced. This could potentially imply a very late neural locus for the judgments, at a level where the stimuli are coded in terms of their perceptual dimensions. To examine this, we used previous measures of hue scaling functions, in which observers rate the proportion of red vs. green or blue vs. yellow present in the

hue. Malkoc, Kay, and Webster (2005) sampled these at 15 deg intervals in the LvsM and SvsLM plane, at a contrast similar to the one we used in the current study. We used their average reported function (Figure 3 of their paper) and then used a spline interpolation to estimate the values for intermediate angles. We then calculated the cone-opponent angles that would be required so that the average of the perceived hues of a pair equaled the chosen reference hue. For example, if one spot appeared 30% blue and 70% green, then the second spot would need to appear 70% blue and 30% green to form a balanced blue-green. These predictions are shown in Figure 4b. Note that the hue angles based on the hue scaling are in good agreement with the average hue loci for the uniform arrays measured in the current study. However, these again fail to predict the mixtures, and in particular, in this case predict large biases in the red averages while very little bias for blue-green, the opposite of the pattern we actually observed. Intuitively, these predictions reflect the fact that, as noted on the figure, unique yellow and blue are very different angular distances from unique red (~60 vs. 135 deg respectively). Thus a pair of hues symmetric about the red locus will have a stronger yellow than blue component, and the pair must be biased toward blue to compensate for this. Conversely, the loci of blue and green are at more similar distances from the blue-green boundary (~40 vs. 35 deg away from a blue-green boundary at 180 deg), so that pairs straddling the blue-green locus should remain perceptually balanced.

Finally, because the stimulus is composed of spatially interleaved elements, as a fourth possibility we considered whether the biases in the mean

for BG could result from chromatic induction between the components. To account for the BG shifts this would need to be asymmetric and stronger from the blue (i.e. so that blue components caused larger hue shifts in the green elements than vice versa). However, this should again predict that the biases should systematically vary with the spatial frequency of the stimuli and would again be expected to show similar shifts for the unique red settings, if the induction is controlled by separable signals along the cone-opponent axes ("Poster child," 2005). Moreover, this account is again based on the assumption that observers base their mean percepts on the phenomenal appearance of the component colors. Induced changes in appearance would therefore need to effectively cancel the biases predicted by the perceived color of the isolated components in the red settings, requiring that the induced bias produce shifts in opposite directions for the BG and R averages (Figure 4b). For these reasons, induction also fails to account for the observed pattern. In summary, whatever the basis for the averaging, these analyses surprisingly suggest that it is not based on simply averaging either the perceived hue of the individual components or on the mean of separable signals along the LvsM and SvsLM cone-opponent axes.

Hue Pairs With A Fixed Component

In the second set of experiments, we modified the procedure so that the hue angle of one of the two components remained fixed, while the observer varied the angle of the second component to set a specified average color. Representative results for 4 observers are again shown in Figure 5, which this

time plots the chosen variable hue as a function of the fixed hue angle. For the blue-green settings, the fixed component was set to angles ranging from 120 to 210 in 10 deg increments. If the average depended on the linear mean of the cone-opponent signals, then the settings should lie along a straight line passing through the blue-green locus (i.e. if the fixed hue were 20 deg off in the green direction then the variable should be set to be 20 deg in the blue direction, etc.) Instead, the matches are clearly curved, again consistent with a weaker response to blue components than the green. In fact as the figure shows, the blue-green settings are again reasonably approximated by assuming the average is based on the log of the S cone signals. However, as we found before, the settings for the remaining three target hues were instead roughly linear. Thus the pattern of results remained similar across the two paradigms.

We chose the fixed-hue conditions because they allowed us to test for possible categorical biases in the matches. These are predicted to occur if the subject represents the colors discretely in terms of their dominant hue (e.g. treating all hues composed of more than 50% blue as simply “blue”). In the extreme, this would mean that whenever the fixed hue lay on the green side of the blue-green boundary, balancing it would only require setting the variable hue to be just on the blue side of the boundary, so that the choices for all fixed hues would remain close to the boundary (or random within the blue category in the limit of complete categorical coding). This extreme is clearly unrealistic, because it also predicts that we would not be able to discriminate between different shades of blue. However, it remains possible that there are categorical

biases that might lead the observer to give extra weight to the dominant hue in the stimulus. To test for these, we followed Webster and Kay (Webster & Kay, 2012) in assuming that the response to the stimulus was composed of two parts, an analog response that varied continuously with hue angle and a discrete response that represented the color category. The response for each color in the pair could then be modeled as a weighted average of these two components, with the balance matched by equating the total responses:

$$(1-\alpha)\Theta_1+\alpha\Theta_{C1} = (1-\alpha)\Theta_2+\alpha\Theta_{C2} \quad (3)$$

Where Θ_1 and Θ_2 are the analog hue responses, Θ_{C1} and Θ_{C2} are the discrete categorical responses to each hue, and α is the weight given to the categorical response. Figure 6 illustrates the matches for different degrees of categorical bias, assuming the matches follow the linear average of the hue angles. These lie along the negative diagonal if there is no categorical bias, and along the fixed blue-green boundary as the limit of complete categorical coding is approached. As we have seen, the blue-green settings themselves are not linear. Yet regardless of their form, an analog representation of the hues predicts that the same matches will occur whether the green or blue component is fixed. In contrast, any categorical bias should instead lead to an asymmetry in the settings, since the variable angle will be pulled toward the category boundary. That is, a weak blue will be sufficient to match a strong fixed green, yet a weak green will suffice when a strong blue is fixed. In the settings, this asymmetry can be revealed by reflecting the matches about the blue-green boundary – exchanging the x and y values in the plots. Without a categorical

effect the two curves should superimpose, while a categorical bias should instead result in a steeper slope for the reflected settings. For the first 3 observers in Figure 5, there is little evidence of a systematic asymmetry between their original and reflected settings. Intriguingly however, the fourth plot shows one observer who shows what appears to be a substantial categorical bias – her reflected settings are consistently steeper. In fact, her settings are reasonably fit by assuming a log compressive S cone response combined with an equally weighted ($\alpha = 0.5$) categorical response.

To formally test for categorical effects, for each fixed hue (e.g. 120deg) we compared the chosen setting for that hue (e.g. 200deg), to the setting predicted when the variable hue was instead assumed to be fixed (i.e. start instead with a fixed hue of 20 deg and see if the variable hue was set to 120deg). Again this should reproduce the original fixed hue if the settings are symmetrical, whereas categorical effects should bias the reciprocal setting toward the reference boundary. To estimate the reciprocal matches, the curves shown in Figure 5 were again interpolated with a cubic spline. Figure 7 plots the settings for each of the fixed hue angles vs. their complements. The 4 panels show the results pooled across observers for each reference hue. These settings are variable, for as seen above the judgments of the average color of the pair is itself noisy. However, for each color the settings do not deviate from the trend predicted by a pure analog response to the colors. Thus, at least for the average settings across all observers, the results argue against a strong presence of a categorical bias in the judgments.

Discussion

While observers were asked to set the pair of colors to a specified “average” hue, our results do not demonstrate or require that they directly experienced this average in the stimulus and thus it would not have made a difference if the array only consisted of two color patches or an 11 by 11 display. For an attribute like color the mean could be directly encoded simply by including mechanisms with large receptive fields. In fact something like this happens in color assimilation, in which colors within patterns tend to fuse together. These can lead to percepts in which the blended mean and individual components are simultaneously perceived, and one account for this is that the patterns is encoded by mechanisms operating in parallel at multiple spatial scales (De Valois & De Valois, 1980). However, the periodicity in our stimuli was again a very low spatial frequency with clear separation between the clearly distinct elements, and assimilation effects were not visible in the arrays. Thus a physiological basis for the averaging is less certain, and perhaps not necessary. An alternative account of the settings is that observers were comparing the “amounts” of different hues, and adjusting the hue angles so that these amounts appeared equal. The “average” for our target hues should occur when the amounts of the two components were equal, and judging this would not require explicitly representing the mixture. Importantly, in the present conditions these comparisons were in principle sufficient for accomplishing the

task. That is, to judge whether the mixture was a balanced blue-green, observers could compare the relative quantities of “blueness” vs. “greenness” in the two chromaticities (though as we note below these need not correspond to the labeled “blueness” or “greenness” in the stimuli). In further studies we are exploring larger color differences or judgments, in which the target hue cannot be indirectly inferred from the constituent hues. Observers are very poor at these tasks, and this supports an account in which observers are simply comparing the relative magnitude of the two colors rather than directly computing an average.

Whatever its basis, the balance of angles chosen for the hues revealed strong and consistent biases in the blue-green region. As noted, these may in part be due to chromatic aberration, but also point to neural nonlinearities. These in turn cannot be accounted for by the independent responses along the cardinal cone-opponent axes, since the biases are not separable along these axes. Moreover, they were not predictable from the color appearance of the individual components. We do not currently have an explanation for the pronounced shifts in the blue-green settings compared to the more linear matches along the other axes we examined.

Another aim of our work was to test whether the judgments were susceptible to categorical biases. Categorical effects in color perception remain controversial. A number of studies have found that visual search or color discrimination is faster for pairs of stimuli that fall within different rather than the same linguistic categories (Franklin et al., 2008; Gilbert et al., 2006; Roberson

& Davidoff, 2000; Winawer et al., 2007; Witzel & Gegenfurtner, 2011).

However, others have failed to replicate these effects (Brown et al., 2011; Webster & Kay, 2012). Moreover, several authors have noted the possibility that these tasks may reflect linguistic effects on the response rather than the perception of the stimuli (Gilbert et al., 2006; Klemfuss, Prinzmetal, & Ivry, 2012; Pilling, Wiggett, Ozgen, & Davies, 2003; Roberson & Davidoff, 2000; Webster & Kay, 2012). In a recent study we explored categorical effects on color appearance using a perceptual grouping task, specifically to limit the influence of a speeded response on the task (Webster & Kay, 2012). The tendency to group shades of blue and green based on similarity showed very little influence of their relationship to the blue-green boundary. In the present study we have similarly found judging the balance between two colors is not strongly affected by their category, and specifically that observers showed little evidence of over-weighting a hue simply because it fell within a linguistic category. Thus these results again suggest that the actual appearance of color may be impervious to the verbal labels used to classify it. Notably, both the grouping and averaging task made no reference to the hues of the individual components (though we did of course ask observers to adjust the balance in terms of a given color label). In a hue scaling task – where color labels are explicitly used to describe the colors – categorical effects do emerge (Webster & Kay, 2012). As also noted, at least one observer (Figure 4d) did show evidence of a fairly strong categorical bias. Our results do not reveal the basis for these possible individual differences. However, it is tempting to speculate

that this observer was to some extent simply labeling the stimuli by their dominant hue, and basing their average on these labels (i.e. without actually “seeing” the stimuli differently). Our paradigm did not manipulate attention in ways that might more clearly isolate preattentive processes in ensemble coding (“Poster child,” 2005). Yet even when the potential for any type of strategy was available to them, it is notable that most observers showed no evidence that color labeling influenced their judgments.

Categorical effects on color appearance are generally weak. For example, even when they have been found they represent only minor biases in performance. Thus the debate is over whether they show up at all, and not whether they dominate the judgment. Given this, it is not surprising that these effects were largely absent in the present task, especially since the variability in the averaging judgment means that it can only provide a fairly insensitive test for categorical coding. What is much more surprising is that the phenomenal color appearance of the individual stimuli did not even seem relevant to the judgment of the average color. That is, the amount of red vs. green or blue vs. yellow, as reported in tasks such as hue scaling, did not predict how red-green or blue-yellow the ensemble appeared. This dissociation between color appearance and performance has also been reported in a recent study of visual search (Lindsey et al., 2010). The present results suggest that whatever the processes or judgments observers were using to balance the hues, they are largely impervious to how the colors are described, and thus reflect perceptual

computations that not only are not be mediated by language, but more importantly may not be mediated by phenomenal color appearance.

References

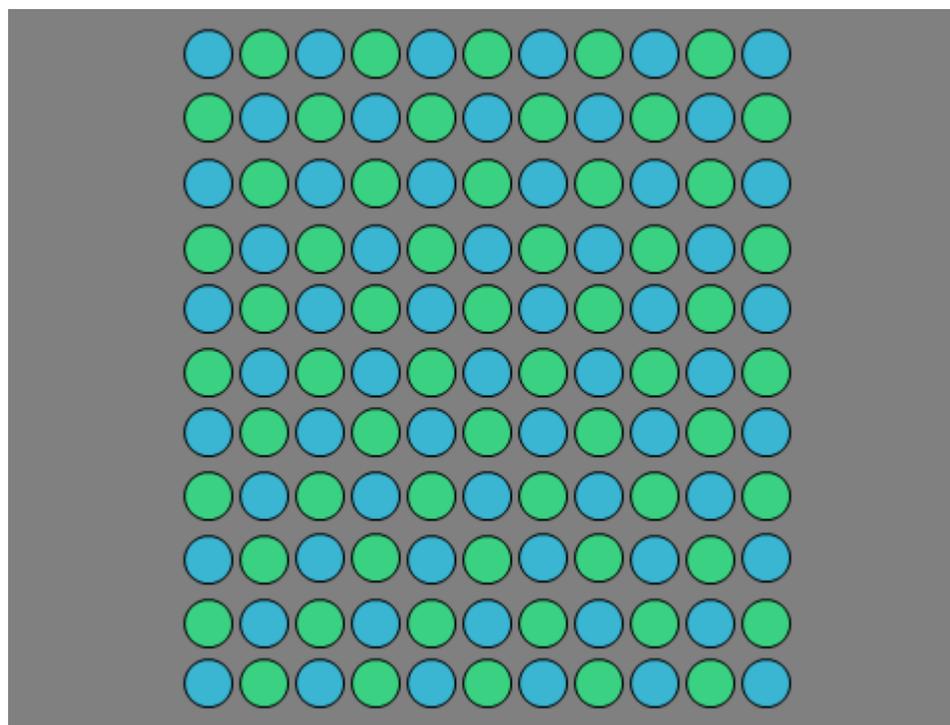
- Alvarez, G. A., & Oliva, A. (2008). The representation of simple ensemble visual features outside the focus of attention. [Research Support, N.I.H., Extramural Research Support, U.S. Gov't, Non-P.H.S.]. *Psychol Sci*, *19*(4), 392-398. doi: 10.1111/j.1467-9280.2008.02098.x
- Ariely, D. (2001). Seeing sets: representation by statistical properties. *Psychol Sci*, *12*(2), 157-162.
- Brainard, D. H., & Wandell, B. A. (1992). Asymmetric color matching: how color appearance depends on the illuminant. *J Opt Soc Am A*, *9*(9), 1433-1448.
- Brown, A. M., Lindsey, D. T., & Guckes, K. M. (2011). Color names, color categories, and color-cued visual search: Sometimes, color perception is not categorical. *Journal of Vision*, *11*, 12-12:11-20. doi: 11.12.2 [pii]10.1167/11.12.2
- Conway, B. R., Chatterjee, S., Field, G. D., Horwitz, G. D., Johnson, E. N., Koida, K., & Mancuso, K. (2010). Advances in color science: from retina to behavior. *J Neurosci*, *30*(45), 14955-14963. doi: 30/45/14955 [pii]10.1523/JNEUROSCI.4348-10.2010
- D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *J Opt Soc Am A*, *3*(10), 1662-1672.
- Dakin, S. C., Bex, P. J., Cass, J. R., & Watt, R. J. (2009). Dissociable effects of attention and crowding on orientation averaging. [Research Support, Non-U.S. Gov't]. *J Vis*, *9*(11), 2821-16. doi: 10.1167/9.11.28
- De Valois, R. L., & De Valois, K. K. (1980). Spatial vision. *Annu Rev Psychol*, *31*, 309-341. doi: 10.1146/annurev.ps.31.020180.001521
- De Valois, R. L., De Valois, K. K., Switkes, E., & Mahon, L. (1997). Hue scaling of isoluminant and cone-specific lights. *Vision Res*, *37*(7), 885-897. doi: S0042698996002349 [pii]
- Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology*, *357*, 241-265.
- Foster, D. H. (2011). Color constancy. *Vision Res*, *51*(7), 674-700. doi: S0042-6989(10)00440-2 [pii]10.1016/j.visres.2010.09.006
- Foster, D. H., & Nascimento, S. M. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proc Biol Sci*, *257*(1349), 115-121. doi: 10.1098/rspb.1994.0103
- Franklin, A., Drivonikou, G. V., Bevis, L., Davies, I. R., Kay, P., & Regier, T. (2008). Categorical perception of color is lateralized to the right hemisphere in infants, but to the left hemisphere in adults. *Proc Natl Acad Sci U S A*, *105*(9), 3221-3225. doi: 0712286105 [pii]10.1073/pnas.0712286105
- Gilbert, A. L., Regier, T., Kay, P., & Ivry, R. B. (2006). Whorf hypothesis is supported in the right visual field but not the left. *Proc Natl Acad Sci U S A*, *103*(2), 489-494. doi: 0509868103 [pii]10.1073/pnas.0509868103
- Goddard, E., Mannion, D. J., McDonald, J. S., Solomon, S. G., & Clifford, C. W. G. . (2010). Combination of subcortical color channels in human visual cortex. *Journal of Vision*, *10*, 1-17.
- Haberman, J., Harp, T., & Whitney, D. (2009). Averaging facial expression over time. *J Vis*, *9*(11). doi: Artn 1 Doi 10.1167/9.11.1

- Haberman, J., & Whitney, D. Efficient summary statistical representation when change localization fails. *Psychon Bull Rev*, 18(5), 855-859. doi: 10.3758/s13423-011-0125-6
- Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current Biology*, 17(17), R751-R753. doi: DOI 10.1016/j.cub.2007.06.039
- Haberman, J., & Whitney, D. (2009). Seeing the Mean: Ensemble Coding for Sets of Faces. *Journal of Experimental Psychology-Human Perception and Performance*, 35(3), 718-734. doi: Doi 10.1037/A0013899
- Howarth, P. A. (1984). The lateral chromatic aberration of the eye. *Ophthalmic Physiol Opt*, 4(3), 223-226.
- Juricevic, I., Land, L., Wilkins, A., & Webster, M. A. (2010). Visual discomfort and natural image statistics. *Perception*, 39(7), 884-899.
- Klemfuss, N., Prinzmetal, W., & Ivry, R. B. (2012). How does language change perception: a cautionary note. *Front Psychol*, 3, 78. doi: 10.3389/fpsyg.2012.00078
- Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research*, 22(9), 1123-1131.
- Krauskopf, J., Williams, D. R., Mandler, M. B., & Brown, A. M. (1986). Higher order color mechanisms. *Vision Res*, 26(1), 23-32. doi: 0042-6989(86)90068-4 [pii]
- Kuehni, R. G. (2004). Variability in unique hue selection: a surprising phenomenon. *Color Research and Application*, 29, 158-162.
- Lennie, P., Krauskopf, J., & Sclar, G. (1990). Chromatic mechanisms in striate cortex of macaque. *J Neurosci*, 10(2), 649-669.
- Lindsey, D. T., Brown, A. M., Reijnen, E., Rich, A. N., Kuzmova, Y. I., & Wolfe, J. M. (2010). Color channels, not color appearance or color categories, guide visual search for desaturated color targets. [Research Support, N.I.H., Extramural Research Support, Non-U.S. Gov't]. *Psychol Sci*, 21(9), 1208-1214. doi: 10.1177/0956797610379861
- MacLeod, D. I., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *J Opt Soc Am*, 69(8), 1183-1186.
- McDermott, J. H., Schemitsch, M., & Simoncelli, E. P. (2013). Summary statistics in auditory perception. *Nature Neuroscience*, 16(4), 493-U169. doi: Doi 10.1038/Nn.3347
- McDermott, K. C., Malkoc, G., Mulligan, J. B., & Webster, M. A. (2010). Adaptation and visual salience. *J Vis*, 10(13), 17. doi: 10.13.17 [pii]10.1167/10.13.17
- Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nat Neurosci*, 4(7), 739-744. doi: 10.1038/8953289532 [pii]
- Piazza, E. A., Sweeny, T. D., Wessel, D., Silver, M. A., & Whitney, D. (2013). Humans use summary statistics to perceive auditory sequences. *Psychol Sci*, 24(8), 1389-1397. doi: 10.1177/0956797612473759
- Pilling, M., Wiggert, A., Ozgen, E., & Davies, I. R. (2003). Is color "categorical perception" really perceptual? *Mem Cognit*, 31(4), 538-551.
- Poirson, A. B., & Wandell, B. A. (1993). Appearance of colored patterns: pattern-color separability. [Research Support, Non-U.S. Gov't Research Support, U.S. Gov't, Non-P.H.S. Research Support, U.S. Gov't, P.H.S.]. *J Opt Soc Am A Opt Image Sci Vis*, 10(12), 2458-2470.
- Poster child. (2005). *Nature Genetics*, 37(5), 447-447.

- Regan, B. C., & Mollon, J.D. (1997). The relative salience of the cardinal axes of colour space in normal and anomalous trichromats In C.R. Cavonius (Ed.), *Colour Vision Deficiencies* (Vol. XIII, pp. 261-270). Kluwer: Dordrecht.
- Roberson, D., & Davidoff, J. (2000). The categorical perception of colors and facial expressions: the effect of verbal interference. *Mem Cognit*, 28(6), 977-986.
- Roberson, D., Pak, H., & Hanley, J. R. (2008). Categorical perception of colour in the left and right visual field is verbally mediated: evidence from Korean. *Cognition*, 107(2), 752-762. doi: S0010-0277(07)00235-1 [pii] 10.1016/j.cognition.2007.09.001
- Robitaille, N., & Harris, I. M. (2011). When more is less: Extraction of summary statistics benefits from larger sets. *J Vis*, 11(12). doi: Artn 18 Doi 10.1167/11.12.18
- Switkes, E. (2008). Contrast salience across three-dimensional chromoluminance space. *Vision Res*, 48(17), 1812-1819. doi: DOI 10.1016/j.visres.2008.05.014
- Watamaniuk, S. N., Sekuler, R., & Williams, D. W. (1989). Direction perception in complex dynamic displays: the integration of direction information. [Research Support, U.S. Gov't, Non-P.H.S.]. *Vision Res*, 29(1), 47-59.
- Webster, M. A., & Kay, P. (2012). Color categories and color appearance. [Research Support, N.I.H., Extramural]. *Cognition*, 122(3), 375-392. doi: 10.1016/j.cognition.2011.11.008
- Webster, M. A., Miyahara, E., Malkoc, G., & Raker, V. E. (2000). Variations in normal color vision. II. Unique hues. *J Opt Soc Am A Opt Image Sci Vis*, 17(9), 1545-1555.
- Webster, M. A., & Mollon, J. D. (1995). Colour constancy influenced by contrast adaptation. *Nature*, 373(6516), 694-698. doi: 10.1038/373694a0
- Webster, M. A., & Wilson, J. A. (2000). Interactions between chromatic adaptation and contrast adaptation in color appearance. *Vision Research*, 40(28), 3801-3816. doi: S0042-6989(00)00238-8 [pii]
- Whitney, D., Haberman, J., & Sweeny, T. (In Press). From textures to crowds: multiple levels of summary statistical perception. . In J.S. Werner & L. M. Chalupa (Eds.), *The New Visual Neurosciences*. Cambridge, Mass.: MIT Press.
- Whitney, D., Haroz, S., & Whitney, D. (2013). Perceiving Group Behavior: Sensitive Ensemble Coding Mechanisms for Biological Motion of Human Crowds. *Journal of Experimental Psychology-Human Perception and Performance*, 39(2), 329-337. doi: Doi 10.1037/A0028712
- Williams, D. W., & Sekuler, R. (1984). Coherent global motion percepts from stochastic local motions. [Research Support, U.S. Gov't, Non-P.H.S.]. *Vision Res*, 24(1), 55-62.
- Winawer, J., Witthoft, N., Frank, M. C., Wu, L., Wade, A. R., & Boroditsky, L. (2007). Russian blues reveal effects of language on color discrimination. *Proc Natl Acad Sci U S A*, 104(19), 7780-7785. doi: 0701644104 [pii] 10.1073/pnas.0701644104
- Witzel, C., & Gegenfurtner, K.R. (2011). Is there a lateralized category effect for color? *Journal of Vision*, 11, 12-16:11-24.
- Wuerger, S. M., Atkinson, P., & Cropper, S. (2005). The cone inputs to the unique-hue mechanisms. *Vision Res*, 45(25-26), 3210-3223. doi: S0042-6989(05)00303-2 [pii] 10.1016/j.visres.2005.06.016
- Zaidi, Q. (1998). Identification of illuminant and object colors: heuristic-based algorithms. *J Opt Soc Am A Opt Image Sci Vis*, 15(7), 1767-1776.

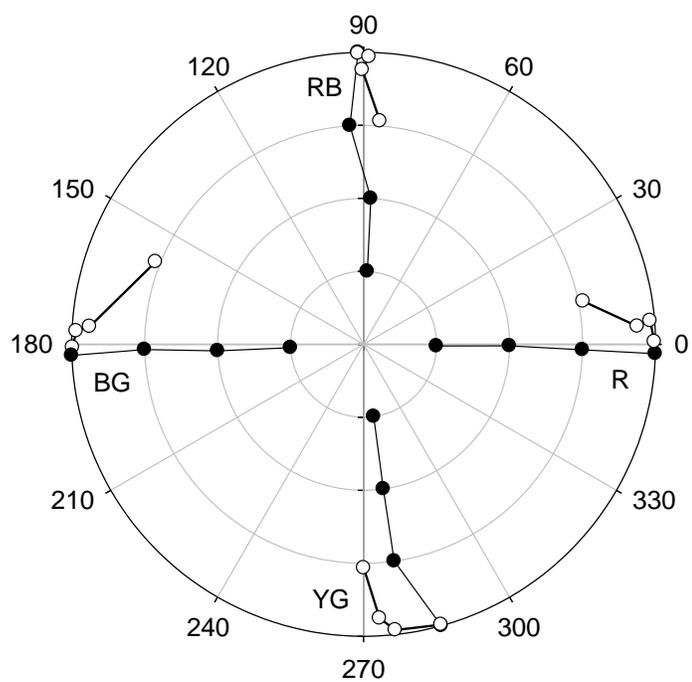
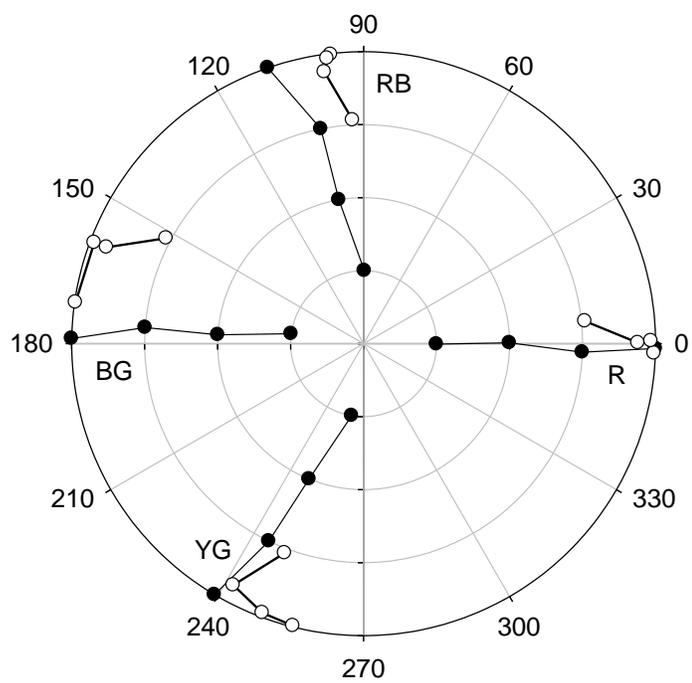
Appendix A

Figure 1. Stimulus Array



Appendix B

Figure 2. Yoked Hue Pairs Settings



Appendix C

Figure 3a. Yoked Hue Pairs Mean Blue-Green Settings

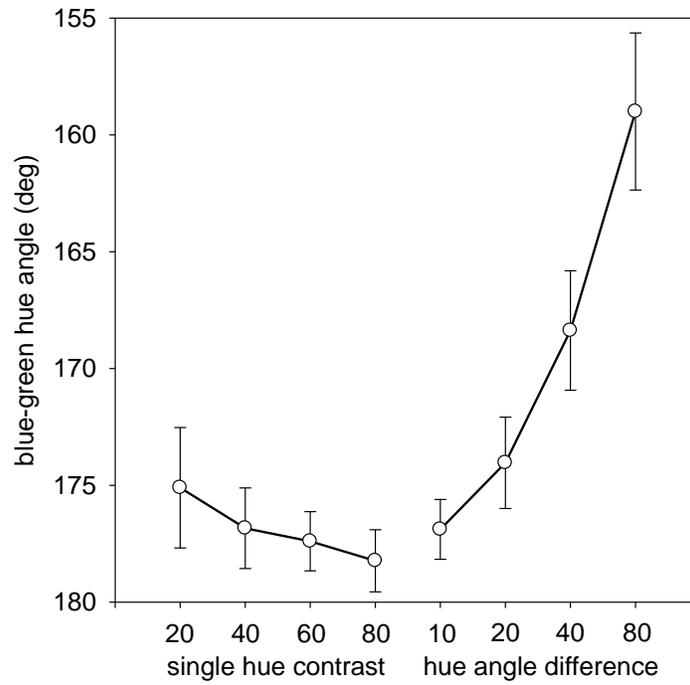
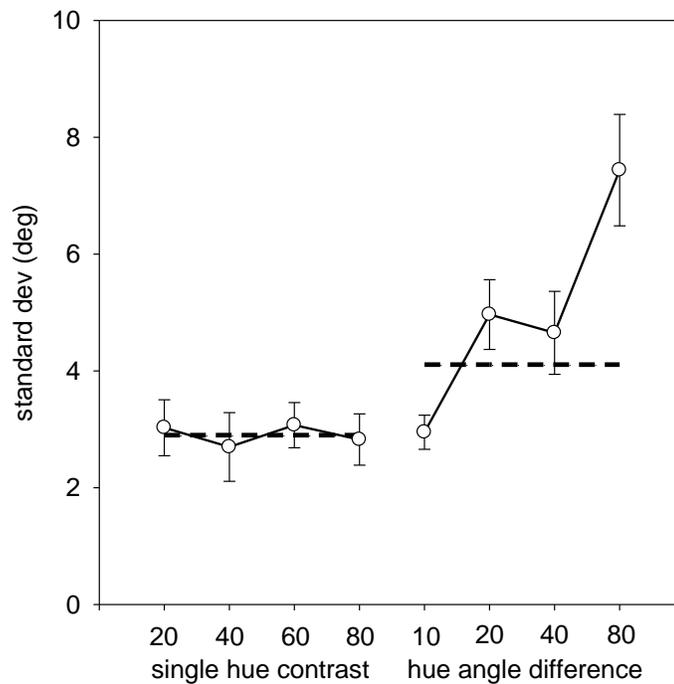
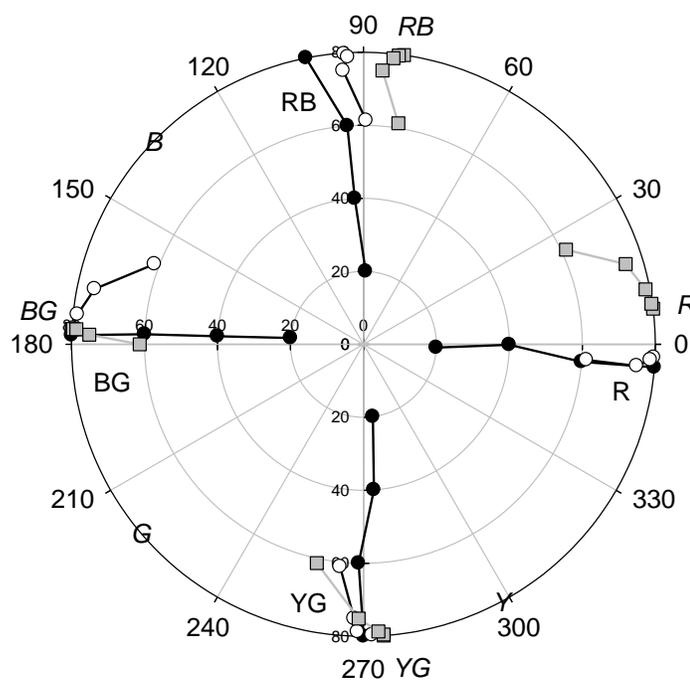
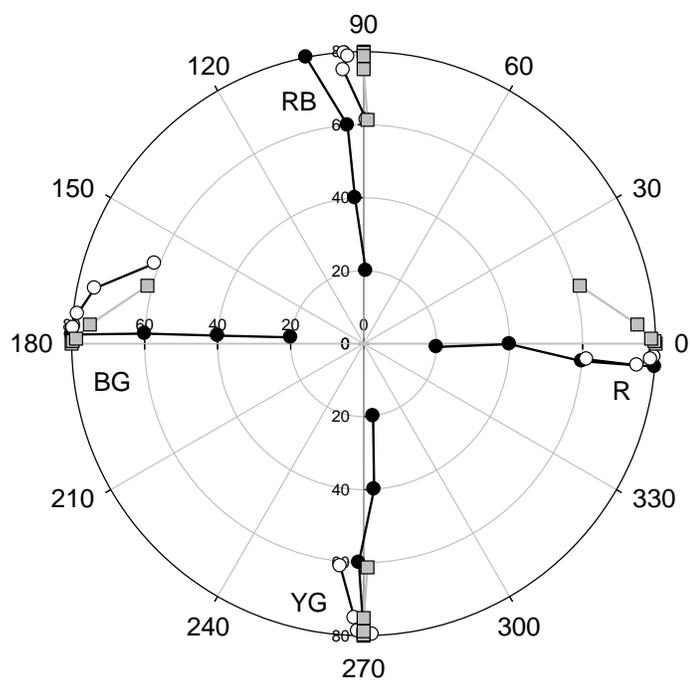


Figure 3b. Yoked Hue Pairs Average Standard Deviation



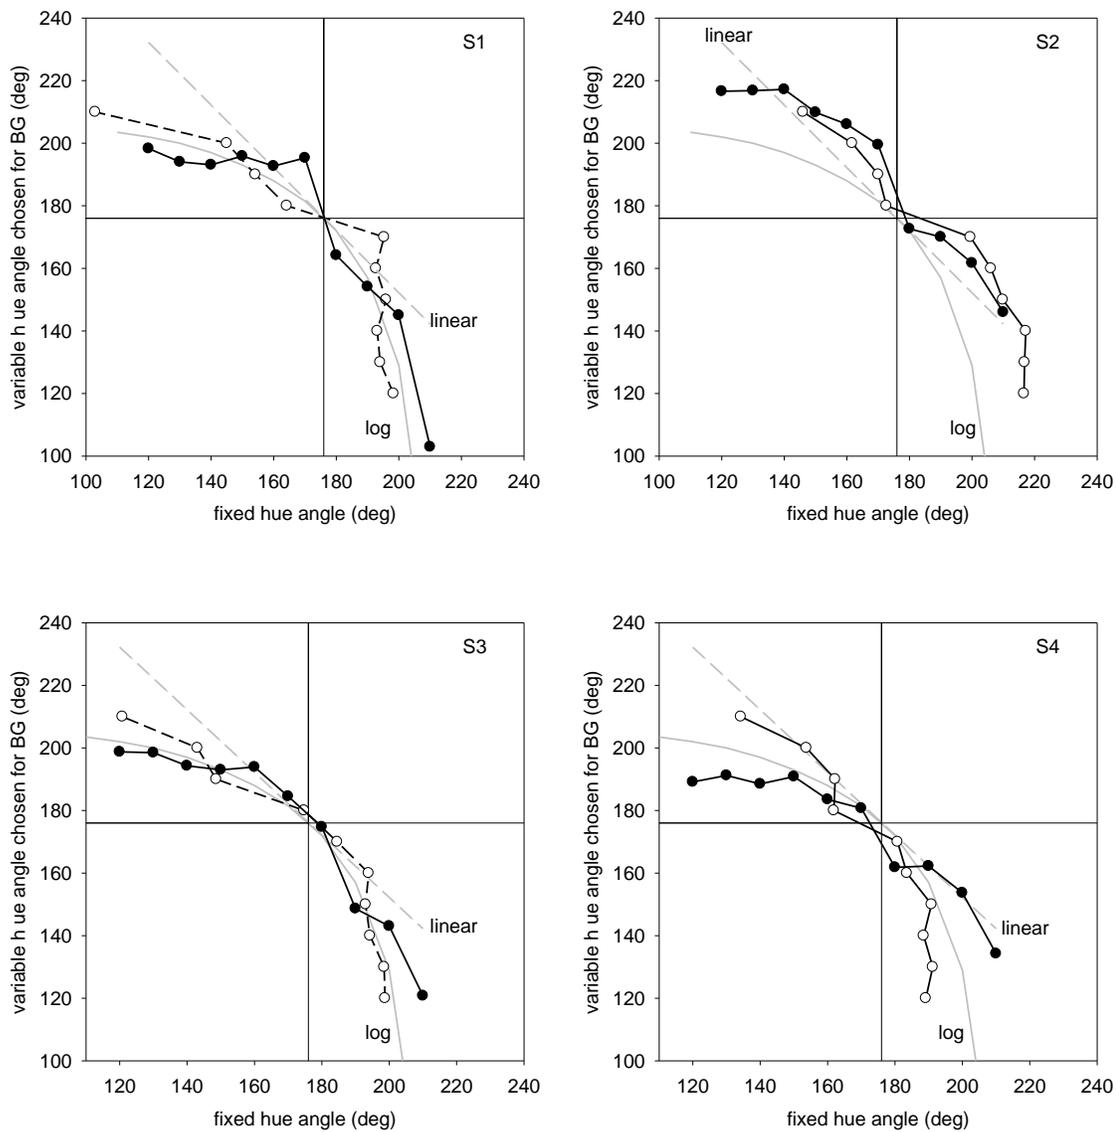
Appendix D

Figure 4. Yoked Hue Pairs Average Settings



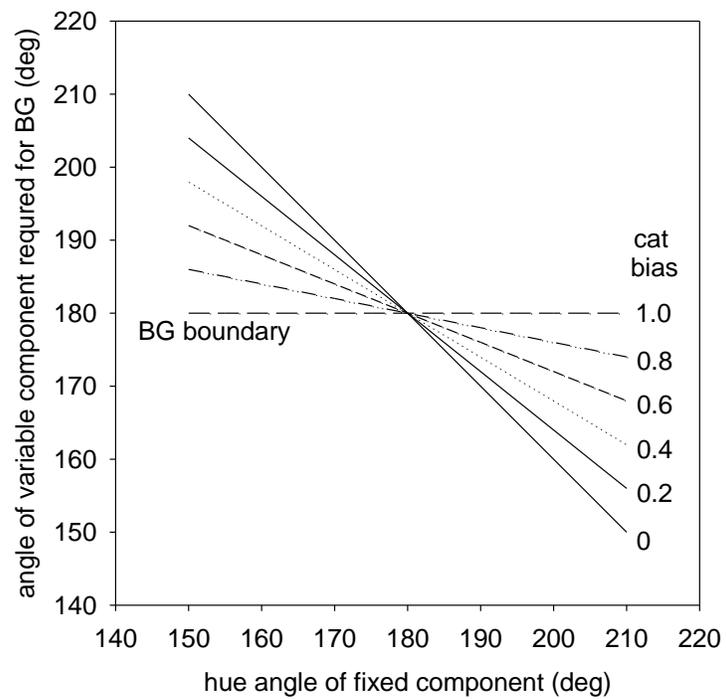
Appendix E

Figure 5. Fixed Single Component Settings



Appendix F

Figure 6. Fixed Single Component Predictions



Appendix G

Figure 7. Fixed Single Component Comparisons

