Color variance and achromatic settings

SIDDHART S. RAJENDRAN AND MICHAEL A. WEBSTER* [ID]

Department of Psychology, University of Nevada Reno, Reno, Nevada 89557, USA
Corresponding author: mwebster@unr.edu

The average color in a scene is a potentially important cue to the illuminant and thus for color constancy, but it remains unknown how well and in what ways observers can estimate the mean chromaticity. We examined this by measuring the variability in "achromatic" settings for stimuli composed of different distributions of colors with varying contrast ranges along the luminance, SvsLM, and LvsM cardinal axes. Observers adjusted the mean chromaticity of the palette to set the average to gray. Variability in the settings increased as chromatic contrast or (to a lesser extent) luminance contrast increased. Signals along the cardinal axes are relatively independent in many detection and discrimination tasks, but showed strong interference in the white estimates. This "cross-masking" and the effects of chromatic variance in general may occur because observers cannot explicitly perceive or represent the mean of a set of qualitatively different hues (e.g., that red and green hues average to gray), and thus may infer the mean only indirectly (e.g., from the relative saturation of different hues).   © 2020 Optical Society of America

https://doi.org/10.1364/JOSAA.382316

1. INTRODUCTION

The stimulus that appears gray or achromatic is fundamental to color vision, for it is the norm relative to which all other colors are encoded [by the direction (hue) or distance (saturation) from gray]. Many studies have examined achromatic settings and how they depend on factors such as visual eccentricity [1,2], the observer's age [3,4], the state of adaptation [5], or the measurement procedures [6]. In general, chromatic sensitivity is highest around the white point relative to other colors, consistent with the predictions for contrast masking in which white corresponds to zero contrast [7–9]. Achromatic settings have also played a central role in color constancy, as a measure of how the perceived colors of objects depend on the spectral characteristics of the illuminant (e.g., [10]). White settings tend to be more variable along bluish–yellowish directions compared to reddish–greenish axes [3,6,11]. This blue–yellow bias has been attributed to weaker sensitivity or greater tolerance for variations in natural daylight spectra, which vary along a blue–yellow axis [12–16].

Most studies of achromatic settings examine how or how well observers can judge the appearance of a single surface. However, natural scenes often include a wide gamut of colors, and can have a strong bias in the average chromaticity [17–21]. It is thus important to understand the extent to which an observer can estimate the average color of such scenes [22]. For example, this would be important for discounting the illuminant to achieve color constancy, assuming a gray world or neutral set of reflectances [23,24]. The ability to extract the mean of a set of

elements has been repeatedly demonstrated in the context of "ensemble coding" or "summary statistics." Ensemble coding is motivated by the observation that the visual system is often insensitive to the individual details comprising a scene, yet highly sensitive to the global statistics of these details such as their mean and variance (see [25] for a review). This finding has been demonstrated for a wide variety of visual attributes including size [26], orientation [27], and motion [28]. It has also been found for high-level features such as facial expressions [29] and direction of gaze [30]. The ability to extract the mean of an ensemble has also been extended to audition, for example, to represent the average of a set of tones [31,32]. The ubiquitous nature of averaging in sensory processing suggests it confers fundamental advantages in perception, and in particular supports efficient representations of image properties [25,33].

A number of studies have examined ensemble coding in color vision (e.g., [34–38]). These have generally shown robust averaging for nearby colors, and weaker sensitivity as the color variance increases, but have not revealed how this averaging is related to basic mechanisms of color coding. Here we systematically probed how averaging and white percepts depend on different cardinal directions in color space [39], and the interactions between these directions. Specifically, our goal was to assess the relative independence of the cardinal chromatic axes in a color averaging task, and explore if the sensitivity of white settings for stimuli having different chromatic contrasts was separable based on these cardinal directions.

2. METHODS

A. Participants

Participants included author SS and 12 students at the University of Nevada, Reno, who were unaware of the specific aims of the study, with different observers participating in different subsets of the experiments. All observers had normal color vision as assessed by the Cambridge Color Test [40] and had normal or corrected-to-normal visual acuity. The studies followed protocols consistent with the Declaration of Helsinki and were approved by University of Nevada, Reno's Institutional Review Board, and participation was with informed consent.

B. Stimuli

Stimuli were displayed on a calibrated NEC color monitor controlled by a Cambridge Research Systems ViSaGe graphics card. The gun luminances were linearized through calibration tables measured with a PhotoResearch PR 655 spectroradiometer. The test stimulus was composed of a 15×15 square array of color patches. Each uniform patch subtended 0.5 deg at the 110 cm viewing distance, and the array was shown centered in the 800 pixels by 600 pixels background of the monitor.

The test array and background (except where noted otherwise) had the same average luminance (30 cd/m^2) and an achromatic chromaticity (CIE 1931 x,y: 0.31, 0.316) equivalent to the chromaticity of Illuminant C. Luminance in the display was defined photometrically and not calibrated for individual observers. The colors of the individual array elements were specified by their LvsM, SvsLM, and luminance contrast relative to the mean, with contrasts scaled according to a version of the Macleod–Boynton [41] (MacLeod and Boynton 1979) color space (see Fig. 1). The scaling was chosen to roughly equate sensitivity for the cardinal axes based on previous studies [42]:

$$\text{LvsM contrast} = (r_{mb} - 0.6568) \times 2754,$$

$$\text{SvsLM contrast} = (b_{mb} - 0.01825) \times 4099,$$

$$\text{LUM contrast} = (\text{Lum} - 30) \times 3.33,$$

where r_{mb} and b_{mb} are the chromaticity coordinates of the stimuli in the MacLeod–Boynton diagram, and the constants 0.6568, 0.01825 are the MacLeod–Boynton coordinates of the background. Lum is the luminance of the element, with the contrasts scaled to range from -100 to 100. For a given contrast level (e.g., 30), the contrasts of the 225 elements were chosen to uniformly and equally sample 225 values along the contrast range (e.g., from -30 to 30), with the spatial locations of the contrast levels randomized. The color of each patch displayed the combination of the randomly assigned contrasts for the three cardinal axes. Different conditions varied in (1) the contrast range of the elements, (2) the chromatic axis used to define the array, or (3) the type of background. The specific conditions are noted for each set of results below.

C. Procedure

Observers viewed the display binocularly in an otherwise dark room and fixation was not controlled. During an experimental
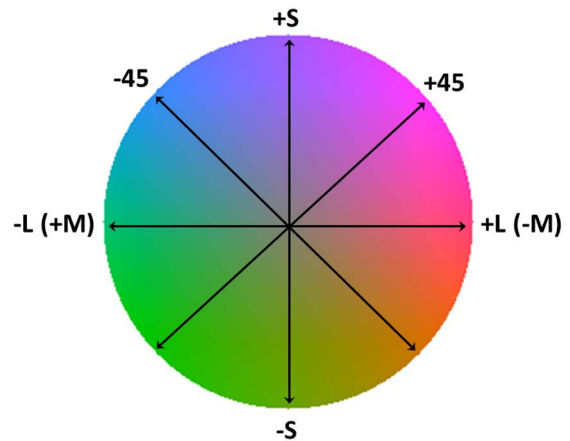
Fig. 1. Axes for chromatic contrast variations in the stimuli. The LvsM and SvsLM axes are the cardinal directions of color space, and correspond to angles of 0–180 and 90–270 deg. The -45 and $+45$ axes correspond to hue angles intermediate to the cardinal axes (based on the scaling used for the space).

session they first adapted to the gray background for 15 s, and then were shown the test palette for 500 ms alternated with 1500 m of the background. The mean of the array was initially set to a random level, with the fixed range of chromaticities varied relative to the mean. For example, if the specified contrast was 30 and the randomly selected mean was 10, then the elements ranged in contrast from -20 to 40. Observers were instructed to adjust the mean chromaticity of the array until the palette appeared unbiased (i.e., with a mean equal to gray), using two pairs of buttons on a handheld keypad. The two pairs separately varied the mean value of the array along two orthogonal axes LvsM and SvsLM or -45 and $+45$. For example, the bottom panel of Fig. 2 shows an array before or after adjusting the palette with a fixed contrast along the LvsM axis and a mean bias in both the LvsM and SvsLM coordinates. The task required adjusting the mean along both axes so that the colors appeared centered on the achromatic point. The mean luminance and luminance contrast remained fixed throughout and could not be adjusted. The stimulus repeated until a setting was selected by a button press, at which point the mean (LvsM and SvsLM, or -45 and $+45$) values were recorded and the program advanced to the next stimulus. During a session participants were shown different color palettes (corresponding to different chromatic axes and to different chromatic and luminance contrasts), each shown 10 times in random order. Two separate sessions were completed for each condition resulting in a total of 20 settings per stimulus for each observer. The standard deviation of these 20 settings was used as a measure of the sensitivity to a color bias in the arrays, and was calculated separately for each adjustment axis (e.g., for the settings along the LvsM or SvsLM axis). The results reported are based on the standard deviations averaged across the observers for each condition. All comparisons are between conditions that were assessed during the same sessions and within the same observers, and are based on repeated-measures ANOVAs comparing the magnitude of the errors across conditions.
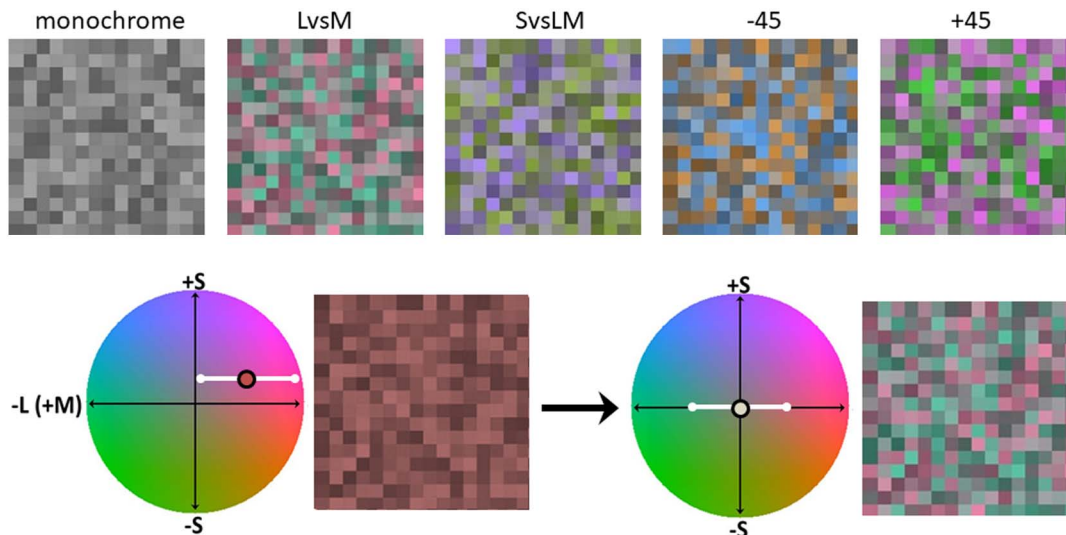
Fig. 2. Top: examples of unbiased stimulus palettes with a fixed luminance contrast and no chromatic contrast (monochrome), or a fixed range of chromatic contrast along the LvsM, SvsLM, −45, or +45 deg axis. In the experiment each array was chromatically biased by adding a random uniform mean chromaticity to the palette, and observers adjusted this mean along the stimulus axis and the orthogonal axis until the palette appeared unbiased. Bottom: example in which the stimulus had a fixed range of chromatic variation along the horizontal LvsM axis (white bar) but with an added random mean shift along both the LvsM and SvsLM axes (reddish purple circle). The mean was adjusted along both the LvsM nd SvsLM axes so that the mean appeared achromatic (gray circle).

3. RESULTS

The general aim of our study was to assess sensitivity to the "white balance" of color arrays and how this depends on the range of color and luminance in the arrays. In all cases we assessed this by asking observers to adjust the mean chromaticity of the palettes so that the mean was achromatic, and used the variance in the repeated estimates as a measure of sensitivity. The stimuli varied in luminance contrast, chromatic contrast, and the axis of the chromatic contrast. In further experiments we also examined the effects of the backgrounds on which the arrays were presented. In the following sections we consider the effects of each of these factors.
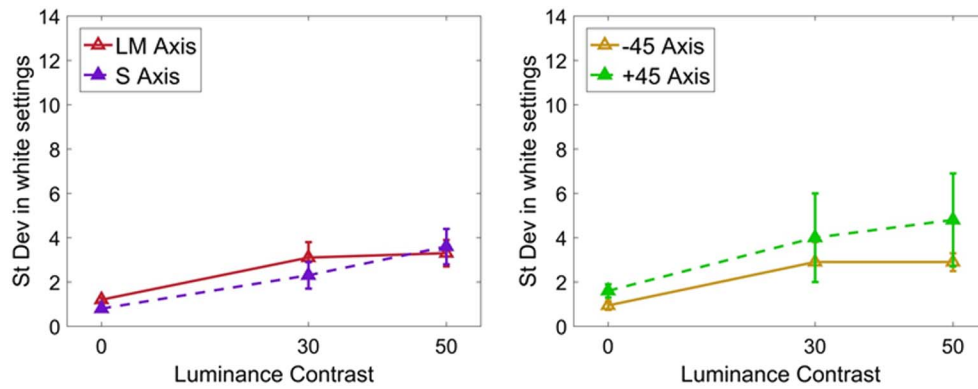
A. Effects of Luminance Contrast

In the first case we examined the achromatic settings in palettes that had uniform (zero contrast) chromaticity, but differed in the level of luminance contrast. The luminance contrast was fixed at 0, 30, or 50, and observers adjusted the chromaticity of the palettes until they appeared gray. Figure 3 shows plots of the resulting achromatic settings, again for palettes that had a range of luminance levels but were a uniform color (i.e., zero chromatic variance). In this and subsequent plots, the symbols show the mean standard deviation of the settings [averaged across the 20 trials and eight observers (left), four observers (right)] as function of the array contrast, and the different lines correspond to the settings along the two chromatic axes. The left panel shows settings when observers adjusted the LvsM and SvsLM levels, while the right panel was for a case (discussed further below) where the adjustment was instead along the −45 and +45 deg axes intermediate to the cardinal axes. In each panel the leftmost points correspond to a uniform square. For this condition observers simply had to adjust the square until it

was indistinguishable from the background, and the variance in the settings is near the nominal thresholds for discriminating a color change along either axis. It should thus be noted that in this case observers were not necessarily judging whether the square appeared achromatic since they may have simply been matching the background. However, a discriminable difference from a gray background at equiluminance does typically appear chromatic at the detection threshold [43,44].

Surprisingly, variance in the achromatic settings was increased when luminance contrast was added to the array. That is, even though the arrays remained chromatically uniform, judging the color was more difficult when the elements randomly varied in luminance. These effects were observed both for the mean settings as well as the settings for individual observers. Thus the increased variance is not simply because observers started to reliably differ in their chosen achromatic points, but because each observer was less reliable in their settings. A two-way ANOVA [luminance contrast (0 versus 30 versus 50) versus response axis (LM versus S)] revealed a main effect of luminance contrast, but no main effect of response axis or a significant interaction between the groups. *Post hoc* comparisons showed that the difference in SD of white settings was between 0 versus 50 contrast levels (p = 0.013). A similar trend is apparent in the blue–yellow settings, though in this case the effect of luminance contrast did not reach significance.

B. Effects of Chromatic Contrast: Cardinal Axes

The next set of conditions varied the chromatic contrast along the LvsM or SvsLM axes, with contrast along one axis fixed at 0, 30, or 50 and the other fixed at 0. Because of the effects of luminance variations noted above, the settings were also made for either 0 or 30 luminance contrast, for a total of 10 unique conditions. Examples of the palettes and the settings are shown

Fig. 3. Effect of luminance contrast on white settings. The standard deviation of white settings along the cardinal (left) or intermediate (right) chromatic axes is plotted as a function of the luminance contrast of the stimulus arrays. Points plot the mean and standard error of the values across observers. Points plot the mean and standard error of the settings across the observers.
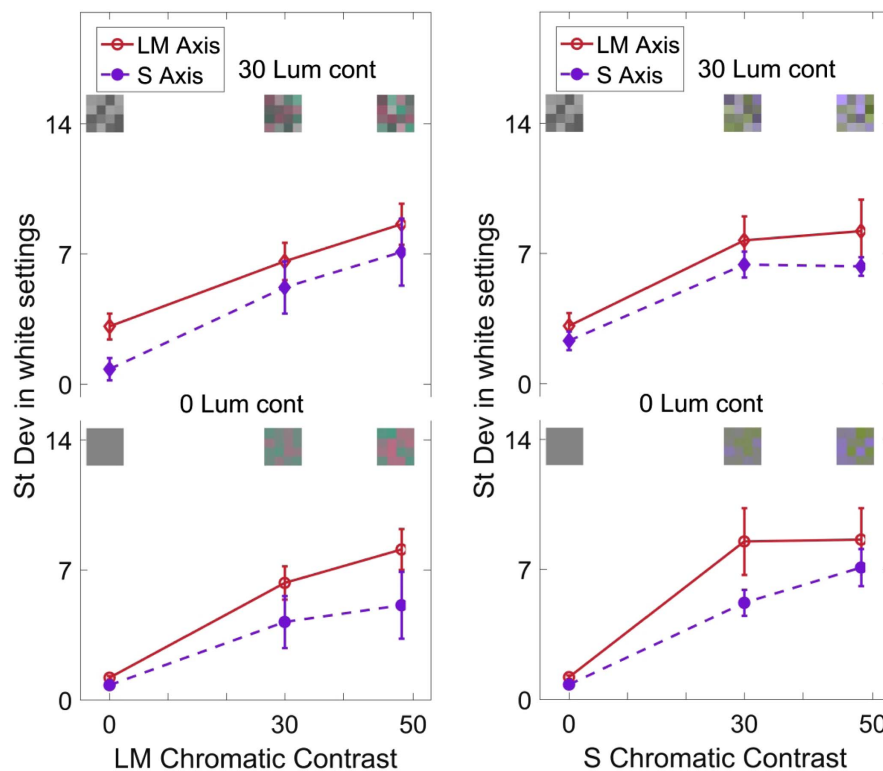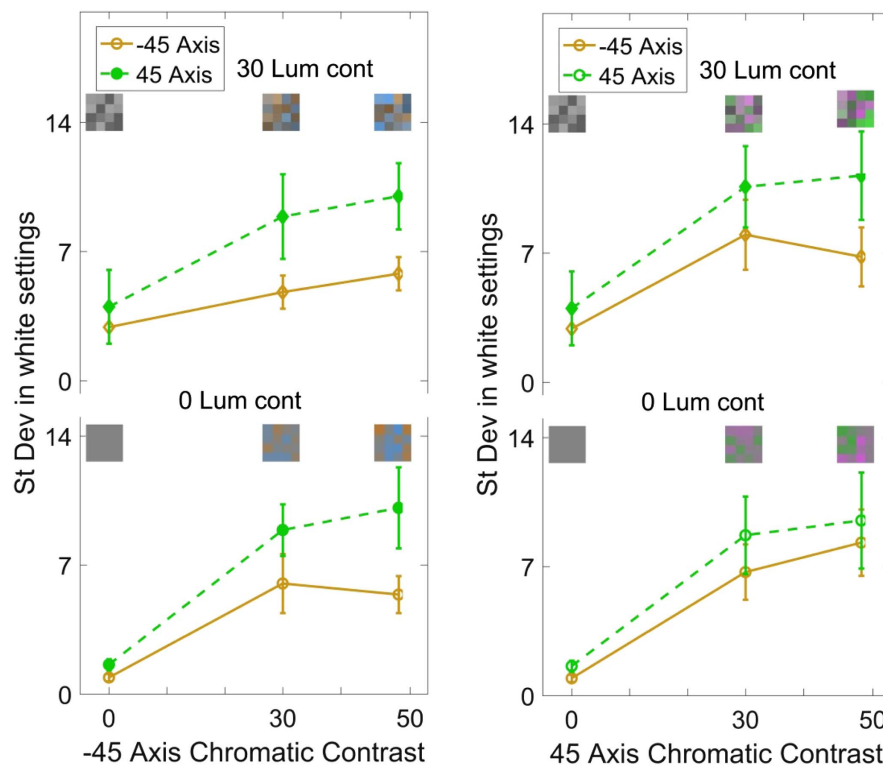

Fig. 4. Standard deviation of the achromatic settings for stimuli with chromatic contrast along the LvsM axis (left) or SvsLM axis (right), and with a luminance contrast of 30 (top) or 0 (bottom). Points plot the mean and standard error across observers for the settings along the LvsM axis (solid lines) or SvsLM axis (dashed lines).

in Fig. 4. The leftmost points replicate the previous results for luminance contrast, with larger standard deviations in the settings when luminance contrast was added to the uniform chromatic array. For both axes, the variance in the settings also increased with increasing chromatic contrast. Importantly, this effect was similar whether the stimulus variance was along or orthogonal to the axis of the judgment. That is, the uncertainty in the LvsM achromatic setting was increased in a similar way whether the stimulus contrast was increased along the LvsM axis or the SvsLM axis. Finally, the results suggest potential differences in the magnitude of the effect along each axis.

However, this difference could reflect a difference in how contrasts along the two axes were scaled.

These effects were formally assessed with a three-way RMANOVA (axis of variance by contrast level by response axis; separately for each luminance contrast), which showed significant main effects for all factors [variance axis, $F(1, 7) = 7.379$, $p = 0.005$; contrast level, $F(2, 14) = 25.076$, $p < 0.001$; response axis, $F(1, 7) = 15.95$, $p = 0.005$]. There was also a significant interaction between the response axis and contrast level [$F(2, 14) = 6.329$, $p = 0.011$]. A *post hoc* analysis of pairwise comparisons between the contrast levels showed that the

Fig. 5. Standard deviation of the achromatic settings for stimuli with chromatic contrast along the −45 deg axis (left) or +45 deg axis (right), and with a luminance contrast of 30 (top) or 0 (bottom). Points plot the mean and standard error across observers for the settings along the −45 deg axis (solid lines) or +45 deg axis (dashed lines).

difference in settings was significant between the 0 versus 30 (p = 0.004), and 0 versus 50 contrast levels (p = 0.003). We also conducted separate three-way RMANOVAs for the LvsM and SvsLM arrays (luminance contrast by chromatic contrast by response axis). For LvsM, there was a main effect only for chromatic contrast [$F(2, 14) = 21.281$, $p < 0.001$]. *Post hoc* pairwise comparisons between the contrast levels showed that the differences between all the chromatic contrast levels were significant (0 versus 30, p = 0.008; 0 versus 50, p = 0.006; 30 versus 50, p = 0.034). Similarly, for the SvsLM arrays showed a main effect of chromatic contrast [$F(2, 14) = 27.18$, $p < 0.001$] with significant differences between the 0 versus 30 (0.002) and 0 versus 50 (0.002) contrast levels.

C. Effects of Chromatic Contrast: Intermediate Chromatic Axes

As noted, a number of previous studies have shown that color discrimination and achromatic settings tend to be less reliable along the "blue–yellow" dimension of color [6,11]. Again, this effect has been attributed to the greater variance in natural lighting along the blue–yellow axis, which might reduce sensitivity to blue–yellow variations. We assessed whether a similar bias would occur for achromatic settings in the variegated color distributions. For this we repeated the same measurements, but for stimuli that were rotated 45 deg in the chromatic plane. One of these axes corresponded to a bluish to yellowish/orange hue variation (thus this axis was close to but did not isolate a pure blue–yellow variation). The second, +45 deg axis corresponded

to hues that varied from reddish–purple to green. Examples of the palettes are shown in Fig. 5. Unlike the preceding conditions, note that any differences in these stimuli cannot result from differences in how the LvsM and SvsLM axes are scaled, for both axes include the same LvsM and SvsLM contrasts, combined in opposite phase.

Here again, variability in the settings increased with increasing chromatic contrast (Fig. 5). Surprisingly, for our conditions the achromatic settings were, if anything, more reliable along the −45 deg (bluish–yellowish) axis when the arrays were monochromatic. Moreover, when the stimulus varied along the +45 deg, settings also tended to be less noisy for this axis. However, the converse was not observed for stimulus variations along the +45 reddish–greenish axis. Thus, in general, our results did not reveal a weaker sensitivity to the −45 deg axis, but again confirmed that achromatic settings became less reliable as the stimulus variance increased. A three-way RMANOVA (stimulus axis by contrast level by response axis; for each luminance contrast) showed that there was a main effect of the response axis [$F(1, 3) = 17.617$, $p = 0.025$] and main effect of the contrast level [$F(2, 6) = 19.088$, $p = 0.003$], but no effect of the stimulus color axis and no interaction between the factors. Separate analyses of the −45 and +45 stimuli (comparing luminance level by chromatic contrast by response axis) showed a significant main effect for chromatic contrast [$F(2, 6) = 14.36$, $p = 0.005$] with the difference between 0 versus 50 being significant (p = 0.038), and no main effect of luminance contrast.
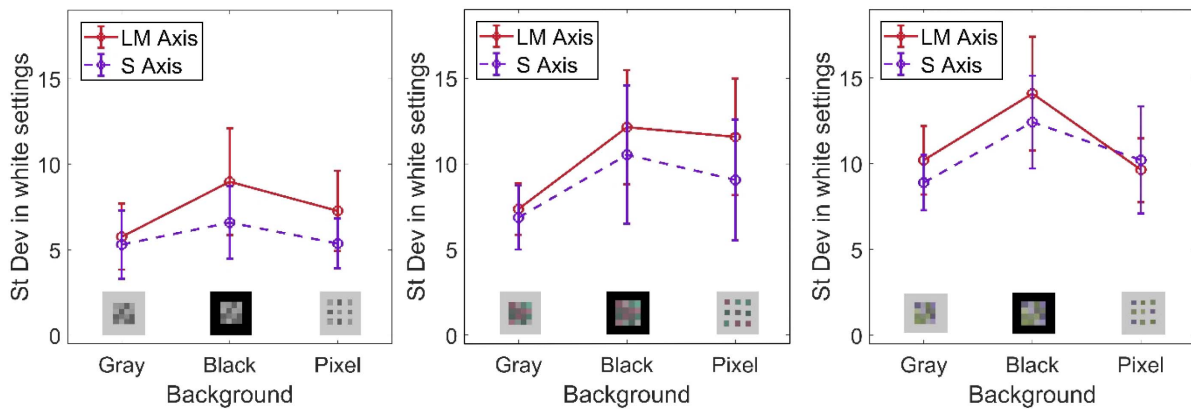
Fig. 6. Effect of the background on the settings. Each panel plots the standard deviation in the achromatic settings for a given stimulus array shown on a gray background, black background, or a gray background interleaved with the array elements. The three panels are for stimulus arrays with luminance contrast but no chromatic contrast (left), luminance contrast and LvM chromatic contrast (middle), or luminance contrast and SvsLM contrast (right). Points plot the mean and standard error across the observers for the adjustments along the LvsM (solid lines) and SvsLM (dashed lines) axes.

For the +45 stimulus dimension the main effects did not reach significance.

D. Effects of the Background

In the final conditions we assessed the effect of the background surrounding the arrays. As noted, in the preceding conditions the background was maintained at the gray average of the arrays, and thus provided a potential reference for setting the mean chromaticity of the array. To remove this cue, we repeated the settings for the cardinal axis contrasts, but on a black background. This increased the variability in the settings, although not significantly (Fig. 6).

As a second manipulation we introduced a gap between each element of the array. The gap subtended 0.5 deg by setting alternate elements in the array to gray. An example of the array is shown in Fig. 7. If the background served as a reference, then the added gap might provide a stronger local cue to the average chromaticity of the array. Alternatively, the gap might *decrease* sensitivity because it removed the potential cue for a mean bias in the color to reflect a uniform illuminant or transparent overlay. Instead, however, this condition did not alter the variability in the white settings. These effects were confirmed using a three-way ANOVA (axis of stimulus variation by background type by response axis), which showed no main effects of the stimulus [$F(2, 10) = 2.18$, $p = .164$], or background [$F(2, 10) = 2.17$, $p = 0.165$], or response axis [$F(1, 5) = 1.8$, $p = 0.238$].

4. DISCUSSION

To summarize, achromatic settings in our stimuli were strongly affected by the variance of the color distribution, becoming more variable as the contrast of the distribution increased. Surprisingly, this occurred even when the contrast in the distribution was along color directions that differed from the chromatic axis along which observers judged the achromatic balance. We consider the implications of these effects for how and how well the visual system might encode and represent the achromatic mean of a color distribution.
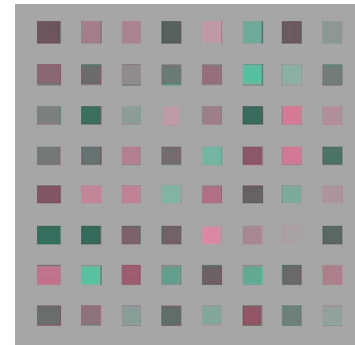
Fig. 7. Example of the stimulus array with gaps between the elements by setting alternate elements to gray.

As noted in the Introduction, the average chromaticity in a scene can provide potential information about the illuminant and thus could provide a powerful cue to color constancy, though several authors have noted that its utility depends in part on an unbiased distribution of reflectances (the "gray world" assumption) [23,24]. Studies of color constancy have demonstrated mechanisms that do respond to mean changes in the stimulus. For example, the cones adapt to the time-averaged signals they are exposed to and thus could adjust to the mean by sampling the image with successive fixations [45,46]. However, this adaptation should actually work against detecting an average color bias since it acts to remove this bias. A uniform shift in the chromaticity (consistent with a global illuminant change) can also be readily distinguished from random changes in the colors within a distribution (consistent with local reflectance changes) supporting relational color constancy [47]. In principle, mechanisms tuned to lower spatial frequencies could also extract the average color of the scene by providing an estimate of the illuminant [48].

As the latter example suggests, one strategy for estimating the mean could be to attend to the low spatial frequency component of the palettes, since this might correspond to the global illumination. There is evidence that the visual system can form layered representations of the illumination and surface color

[49], and transparency effects have been widely studied [50]. Indeed, we specifically considered the possibility that observers might perceive the palettes as unbiased but with an overlying transparent layer of biased color. The control experiment with gaps added between the palette elements was designed to reduce this impression of transparency. However, this did not alter the pattern of settings, arguing against this interpretation. While we did not vary the element size, the perception of the mean also depends strongly on this, since both optical and neural blur would increasingly blend the perceived color as the spatial frequency of the array increases [51,52], ultimately rendering the mean easier to perceive.

Our principal finding is that the extraction of the average is more difficult as the variance in the distribution increases. One mechanism that could give rise to this effect for color is contrast masking. Contrast discrimination thresholds increase with increasing pedestal contrast for both luminance and chromatic contrast, consistent with a saturating nonlinearity in the contrast response [7,8,53,54]. Thus the higher-variance patterns might reduce sensitivity to a chromatic bias. Contrast-dependent losses of sensitivity could also occur with contrast adaptation, which produces stronger losses for higher adapting contrasts [55]. Our stimuli involved brief and intermittent presentations and thus limited but did not completely control for potential adaptation effects.

However, a problem with both masking and adaptation accounts is that both are strongly color-selective. That is, masking or adaptation along the LvsM axis strongly affects sensitivity along the LvsM axis but not along the SvsLM axis or vice versa [7,55–57]. Instead, we found pronounced "cross-axis" interactions on the settings. Specifically, judging the LvsM balance was affected in similar ways whether the color distribution varied only along the LvsM axis or the SvsLM axis, and was also affected even when the distribution varied only in luminance. These cross-axis interactions argue against a simple nonlinearity as the basis for the effects, and suggest they may reflect a fundamentally different cause. Related to this, they may also reflect a fundamentally different judgment. In masking studies the task of the observers is to discriminate between two colors or color distributions, without regard to their color appearance. In contrast, our task required observers to directly judge the appearance.

The focus on appearance rather than discrimination in our task may make the achromatic settings more akin to ensemble coding. As we noted, ensemble coding of the mean has been demonstrated for a wide range of stimuli [25], and like our color settings, becomes less precise with increasing variance in the ensemble [35]. The mechanisms mediating this representation are not well understood. However, our results suggest that the "averaging" does not occur independently within the cardinal mechanisms underlying early post-receptoral color coding, potentially pointing to higher stages of color processing. Moreover, we found little evidence for increased variability along the blue–yellow or daylight locus characteristic of achromatic settings [6,11]. This may be because the noise introduced by the color variations in the palette became more important in limiting performance than the intrinsic sensitivity to white, but at least for our limited settings this suggests that there do not

appear to be any privileged axes for the averaging. How then was an average estimated?

An important issue for understanding the underlying processes is whether the achromatic settings reflected an explicit versus implicit percept of gray in the stimulus—that is whether the visual system forms an actual representation of the mean of the distribution or only infers it indirectly. In the case of color—as well as potentially many of the other stimulus dimensions studied in ensemble coding—this issue remains unresolved. However, it is worth noting that it is not intuitive that the mean of two complementary colors is gray; and in fact, when naïve observers are asked to select the complement of a color so that the mean is gray, their settings are highly variable [58]. Thus, at least anecdotally there seems to be little conscious access to the mean. Moreover, it is possible that observers were not directly sensing the mean hue, but instead gauging the relative contrast of different hues. For example, the white point might have been set by adjusting the reds and greens in the palette to appear equal in saturation. In fact some observers reported this strategy after testing. By this account, increasing the variance along other color directions might make these relative saturation judgments more difficult, because they introduced random hue variations in the stimuli to be compared.

Finally, we note that our settings were confined to the perception of unbiased color distributions, and very different factors may come into play when the average color of the stimulus strongly deviates from white. For example, a number of studies have found that in a variegated distribution, the more saturated colors tend to dominate the average perceived color [59–61], while lightness judgments are biased by the tendency to fixate the lightest parts of objects [62]. This makes sense because in scenes with lighting-induced variation the purest colors or brightest elements may be interpreted as the most undiluted sample of the surface color. However, such effects again point to a role for higher-level inferences in the perception of the average color.

REFERENCES

1. D. Beer, J. Wortman, G. Horwitz, and D. MacLeod, "Compensation of white for macular filtering [abstract]," J. Vis. **5**(8):282, 282 (2005).
2. M. A. Webster and D. Leonard, "Adaptation and perceptual norms in color vision," J. Opt. Soc. Am. A **25**, 2817–2825 (2008).
3. J. S. Werner and B. E. Schefrin, "Loci of achromatic points throughout the life span," J. Opt. Soc. Am. A **10**, 1509–1516 (1993).
4. S. M. Wuerger, "Colour constancy across the life span: Evidence for compensatory mechanisms," PLoS One **8**, e63921 (2013).
5. P. B. Delahunt, M. A. Webster, L. Ma, and J. S. Werner, "Long-term renormalization of chromatic mechanisms following cataract surgery," Vis. Neurosci. **21**, 301–307 (2004).
6. T. Chauhan, E. Perales, K. Xiao, E. Hird, D. Karatzas, and S. Wuerger, "The achromatic locus: effect of navigation direction in color space," J. Vis. **14**(1):25, 25 (2014).
7. J. Krauskopf and K. Gegenfurtner, "Color discrimination and adaptation," Vis. Res. **32**, 2165–2175 (1992).
8. E. Switkes, A. Bradley, and K. K. De Valois, "Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings," J. Opt. Soc. Am. A **5**, 1149–1162 (1988).

9. D. I. A. MacLeod, "Colour discrimination, colour constancy, and natural scene statistics (The Verriest Lecture)," in *Normal and Defective Colour Vision*, J. D. Mollon, J. Pokorny, and K. Knoblauch, eds. (Oxford University, 2003), pp. 189–217.

10. D. H. Brainard and B. A. Wandell, "Asymmetric color matching: how color appearance depends on the illuminant," J. Opt. Soc. Am. A **9**, 1433–1448 (1992).

11. J. M. Bosten, R. D. Beer, and D. I. MacLeod, "What is white?" J. Vis. **15**(16):5, 5 (2015).

12. B. Pearce, S. Crichton, M. Mackiewicz, G. D. Finlayson, and A. Hurlbert, "Chromatic illumination discrimination ability reveals that human colour constancy is optimised for blue daylight illuminations," PLoS One **9**, e87989 (2014).

13. K. C. McDermott and M. A. Webster, "Uniform color spaces and natural image statistics," J. Opt. Soc. Am. A **29**, A182–A187 (2012).

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

15. A. D. Winkler, L. Spillmann, J. S. Werner, and M. A. Webster, "Asymmetries in blue-yellow color perception and in the color of 'the dress'," Curr. Biol. **25**, R547–R548 (2015).

16. D. Weiss, C. Witzel, and K. Gegenfurtner, "Determinants of colour constancy and the blue bias," Perception **8**, 2041669517739635 (2017).

17. M. A. Webster and J. D. Mollon, "Adaptation and the color statistics of natural images," Vis. Res. **37**, 3283–3298 (1997).

18. G. J. Burton and I. R. Moorhead, "Color and spatial structure in natural scenes," Appl. Opt. **26**, 157–170 (1987).

19. D. L. Ruderman, T. W. Cronin, and C. C. Chiao, "Statistics of cone responses to natural images: implications for visual coding," J. Opt. Soc. Am. A **15**, 2036–2045 (1998).

20. J. M. Linhares, P. D. Pinto, and S. M. Nascimento, "The number of discernible colors in natural scenes," J. Opt. Soc. Am. A **25**, 2918–2924 (2008).

21. M. A. Webster, Y. Mizokami, and S. M. Webster, "Seasonal variations in the color statistics of natural images," Network **18**, 213–233 (2007).

22. S. M. Nascimento, V. M. de Almeida, P. T. Fiadeiro, and D. H. Foster, "Effect of scene complexity on colour constancy with real three-dimensional scenes and objects," Perception **34**, 947–950 (2005).

23. H. E. Smithson, "Sensory, computational and cognitive components of human colour constancy," Philos. Trans. R. Soc. London B **360**, 1329–1346 (2005).

24. D. H. Foster, "Color constancy," Vis. Res. **51**, 674–700 (2011).

25. D. Whitney and A. Y. Leib, "Ensemble perception," Annu. Rev. Psychology **69**, 105–129 (2018).

26. D. Ariely, "Seeing sets: representation by statistical properties," Psychol. Sci. **12**, 157–162 (2001).

27. N. Robitaille and I. M. Harris, "When more is less: extraction of summary statistics benefits from larger sets," J. Vis. **11**(12):18, 18 (2011).

28. S. N. Watamaniuk and A. Duchon, "The human visual system averages speed information," Vis. Res. **32**, 931–941 (1992).

29. J. Haberman and D. Whitney, "Seeing the mean: ensemble coding for sets of faces," J. Exp. Psychol. **35**, 718–734 (2009).

30. T. D. Sweeny and D. Whitney, "Perceiving crowd attention: ensemble perception of a crowd's gaze," Psychol. Sci. **25**, 1903–1913 (2014).

31. J. H. McDermott, M. Schemitsch, and E. P. Simoncelli, "Summary statistics in auditory perception," Nat. Neurosci. **16**, 493–498 (2013).

32. E. A. Piazza, T. D. Sweeny, D. Wessel, M. A. Silver, and D. Whitney, "Humans use summary statistics to perceive auditory sequences," Psychol. Sci. **24**, 1389–1397 (2013).

33. G. A. Alvarez, "Representing multiple objects as an ensemble enhances visual cognition," Trends Cogn. Sci. **15**, 122–131 (2011).

34. J. Maule and A. Franklin, "Accurate rapid averaging of multihue ensembles is due to a limited capacity subsampling mechanism," J. Opt. Soc. Am. A **33**, A22–A29 (2016).

35. J. Maule and A. Franklin, "Effects of ensemble complexity and perceptual similarity on rapid averaging of hue," J. Vis. **15**(4):6, 6 (2015).

36. J. Maule, C. Witzel, and A. Franklin, "Getting the gist of multiple hues: metric and categorical effects on ensemble perception of hue," J. Opt. Soc. Am. A **31**, A93–A102 (2014).

37. J. Webster, P. Kay, and M. A. Webster, "Perceiving the average hue of color arrays," J. Opt. Soc. Am. A **31**, A283–A292 (2014).

38. A. Chetverikov, G. Campana, and A. Kristjansson, "Representing color ensembles," Psychol. Sci. **28**, 1510–1517 (2017).

39. J. Krauskopf, D. R. Williams, and D. W. Heeley, "Cardinal directions of color space," Vis. Res. **22**, 1123–1131 (1982).

40. B. C. Regan, J. P. Reffin, and J. D. Mollon, "Luminance noise and the rapid determination of discrimination ellipses in colour deficiency," Vis. Res. **34**, 1279–1299 (1994).

41. D. I. MacLeod and R. M. Boynton, "Chromaticity diagram showing cone excitation by stimuli of equal luminance," J. Opt. Soc. Am. **69**, 1183–1186 (1979).

42. M. A. Webster, E. Miyahara, G. Malkoc, and V. E. Raker, "Variations in normal color vision. II. Unique hues," J. Opt. Soc. Am. A **17**, 1545–1555 (2000).

43. J. Krauskopf, D. R. Williams, M. B. Mandler, and A. M. Brown, "Higher order color mechanisms," Vis. Res. **26**, 23–32 (1986).

44. M. A. Webster, K. K. De Valois, and E. Switkes, "Orientation and spatial-frequency discrimination for luminance and chromatic gratings," J. Opt. Soc. Am. A **7**, 1034–1049 (1990).

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

46. M. D'Zmura and P. Lennie, "Mechanisms of color constancy," J. Opt. Soc. Am. A **3**, 1662–1672 (1986).

47. D. H. Foster and S. M. Nascimento, "Relational colour constancy from invariant cone-excitation ratios," Proc. R. Soc. B **257**, 115–121 (1994).

48. E. L. Dixon and A. G. Shapiro, "Spatial filtering, color constancy, and the color-changing dress," J. Vis. **17**(3):7, 7 (2017).

49. B. L. Anderson and J. Winawer, "Image segmentation and lightness perception," Nature **434**, 79–83 (2005).

50. B. L. Anderson, "The perceptual representation of transparency, lightness, and gloss," in *Oxford Handbook of Perceptual Organization*, J. Wagemans, ed. (Oxford University, 2015), pp. 466–483.

51. A. B. Poirson and B. A. Wandell, "Pattern-color separable pathways predict sensitivity to simple colored patterns," Vis. Res. **36**, 515–526 (1996).

52. K. T. Mullen, "The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings," J. Physiol. **359**, 381–400 (1985).

53. G. E. Legge and J. M. Foley, "Contrast masking in human vision," J. Opt. Soc. Am. **70**, 1458–1471 (1980).

54. C. Chen, J. M. Foley, and D. H. Brainard, "Detection of chromoluminance patterns on chromoluminance pedestals I: threshold measurements," Vis. Res. **40**, 773–788 (2000).

55. M. A. Webster and J. D. Mollon, "The influence of contrast adaptation on color appearance," Vis. Res. **34**, 1993–2020 (1994).

56. A. Li and P. Lennie, "Mechanisms underlying segmentation of colored textures," Vis. Res. **37**, 83–97 (1997).

57. M. J. Sankeralli and K. T. Mullen, "Postreceptoral chromatic detection mechanisms revealed by noise masking in three-dimensional cone contrast space," J. Opt. Soc. Am. A **14**, 2633–2646 (1997).

58. M. A. Webster, "The Verriest Lecture: Adventures in blue and yellow," J. Opt. Soc. Am. A **37**, V1–V14 (1997).

59. I. Kuriki, "Testing the possibility of average-color perception from multi-colored patterns," Opt. Rev. **11**, 249–257 (2004).

60. E. Kimura, "Averaging colors of multicolor mosaics," J. Opt. Soc. Am. A **35**, B43–B54 (2018).

61. D. Wollschlager and B. L. Anderson, "The role of layered scene representations in color appearance," Current biology **19**, 430–435 (2009).

62. M. Toscani, M. Valsecchi, and K. R. Gegenfurtner, "Optimal sampling of visual information for lightness judgments," Proc. Natl. Acad. Sci. U. S. A. **110**, 11163–11168 (2013).