PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES

Fundamentally different representations of color and motion revealed by individual differences in perceptual scaling

Kara J. Emery[a,b,1], Vicki J. Volbrecht[c], David H. Peterzell[d,e], and Michael A. Webster[b,1]

Edited by Thomas Albright, Salk Institute for Biological Studies, La Jolla, CA; received February 7, 2022; accepted November 1, 2022

The coordinate frames for color and motion are often defined by three dimensions (e.g., responses from the three types of human cone photoreceptors for color and the three dimensions of space for motion). Does this common dimensionality lead to similar perceptual representations? Here we show that the organizational principles for the representation of hue and motion direction are instead profoundly different. We compared observers' judgments of hue and motion direction using functionally equivalent stimulus metrics, behavioral tasks, and computational analyses, and used the pattern of individual differences to decode the underlying representational structure for these features. Hue judgments were assessed using a standard "hue-scaling" task (i.e., judging the proportion of red/green and blue/yellow in each hue). Motion judgments were measured using a "motion-scaling" task (i.e., judging the proportion of left/right and up/down motion in moving dots). Analyses of the interobserver variability in hue scaling revealed multiple independent factors limited to different local regions of color space. This is inconsistent with the influences across a broad range of hues predicted by conventional color-opponent models. In contrast, variations in motion scaling were characterized by more global factors plausibly related to variation in the relative weightings of the cardinal spatial axes. These results suggest that although the coordinate frames for specifying color and motion share a common dimensional structure, the perceptual coding principles for hue and motion direction are distinct. These differences might reflect a distinction between the computational strategies required for the visual analysis of spatial vs. nonspatial attributes of the world.

visual perception | sensory coding | perceptual representation

Significance

The visual attributes of color and motion have often been compared to identify different processing streams within the visual system. Here, we show that the judgments people make about hue and motion direction – in similar tasks – also reveal very different representational principles. Individual differences in perceived motion direction reflect the structure of visual space (e.g., varying in the relative weighting of horizontal and vertical), while variations in hue percepts do not reveal privileged chromatic axes or a global coordinate frame. This distinction may reflect fundamental differences in how information about hue and motion direction are used for vision.

Normal human color vision is trichromatic. Any spectrum can be matched by varying the intensities of three suitable primary lights, and this is because the spectra are sampled by three independent classes of cones. These cone types differ based on whether their peak sensitivity is to short (S), medium (M), or long (L) wavelengths. The response in any given cone is univariant (i.e., encoding only the total number of quanta absorbed and not their wavelength). Thus, the triplet of cone excitations initially limits the visual representation of the spectral distribution of light to three dimensions. While trichromacy explains why different spectra match in terms of the three-dimensional set of cone excitations, it makes no predictions about how colors appear. Moreover, color appearance can be affected by many variables, including the surrounding context, and thus cannot be predicted from the local cone excitations alone (1). However, at its simplest level (e.g., a single uniform color on a neutral background), our experience of a color can be defined by the three perceptual attributes of hue, saturation, and brightness. In conventional color-opponent theory, these attributes depend on three perceptual mechanisms that mediate red vs. green, blue vs. yellow, and white vs. black sensations (2, 3). By this account, it is assumed that all hues depend on the relative activity of the underlying red–green and blue–yellow responses.

However, how the cone signals are combined and represented to give rise to color sensations remains poorly understood. Color coding is transformed at many stages along the visual pathway. Cells in the retina and lateral geniculate carry chromatic information within two subsystems that compare the L vs. M cones or the S vs. L + M cones (4, 5). These "cardinal mechanisms" are then combined throughout visual cortex to form multiple "higher order" mechanisms tuned to different chromatic directions that may be further recoded at different stages along the ventral visual stream (6, 7). Ultimately, these more advanced stages of color coding show response patterns that more closely align with perceptual similarities among colors than do earlier stages of coding (e.g., the LGN, V1) (8–13). Throughout these processing stages, there is little evidence for the emergence of the red–green and blue–yellow perceptual dimensions of conventional color-opponent theory. For example, stimuli that isolate responses in either of the cardinal chromatic mechanisms do not correspond to the stimuli that

Author affiliations: [a]Center for Data Science, New York University, New York, NY 10011; [b]Graduate Program in Integrative Neuroscience, University of Nevada, Reno, NV 89557; [c]Department of Psychology, Colorado State University, Fort Collins, CO 80523; [d]School of Psychology, Fielding Graduate University, Santa Barbara, CA 93105; and [e]John F. Kennedy School of Psychology, National University, Pleasant Hill, CA 94523

Author contributions: K.J.E. and M.A.W. designed research; K.J.E. performed research; K.J.E., V.J.V., D.H.P., and M.A.W. contributed new reagents/analytic tools; K.J.E. analyzed data; and K.J.E., V.J.V., D.H.P., and M.A.W. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: kara.emery@nyu.edu or mwebster@unr.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2202262120/-/DCSupplemental.

Published January 20, 2023.

appear pure red–green or blue–yellow (14), and the notion that these hues represent isolation of any mechanism is difficult to reconcile with the distributed population code implied by multiple higher order color mechanisms (15). Moreover, perceptual studies are increasingly questioning the superordinate status of red–green and blue–yellow sensations compared with other hues (16–19). Nevertheless, color-opponent theory and the idea that all hue percepts are represented in terms of an underlying set of unique or primary percepts remains a standard account of color appearance (20).

Previous studies have explored the perceptual representation of hue by analyzing individual differences in hue-scaling judgments (16, 17, 21). Many aspects of perception exhibit large and reliable interobserver variation, and the patterns of such differences can provide powerful clues to the processes underlying sensory coding (22, 23). If hue is constructed from a red–green and blue–yellow scaffold, then differences in hue judgments between observers should reflect variations in these underlying perceptual primaries. For example, under this theory, differences in purple, orange, and red percepts should be partially correlated, because they all include a common underlying red component. Surprisingly, however, measurements of these variations instead show that only hues quite similar in appearance covary (16, 17), suggesting that hue is not coded by the responses of underlying red–green and blue–yellow perceptual mechanisms. In other words, the perceptual representation of hue does not appear to reflect a space defined by a small set of perceptual primaries.

Here we ask whether this surprising pattern of individual differences for hue reveals something unique about how the brain represents chromatic information, or whether it is a general property of sensory representations, or potentially of the tasks used to measure them. To examine this, we determined whether similar representations are revealed when individuals judge different visual attributes that are constrained (albeit in different ways) to share a common dimensional structure. Specifically, we compared judgments of hue (which varies with the direction or angle in color space, relative to the achromatic point) with judgments of motion direction (of different angles in visual space, relative to a static stimulus). While chromatic and motion information interact at many levels of the visual system, motion and color have provided some of the strongest evidence for separate visual streams, with the pathways carrying these features coursing in parallel through the early visual system, and with motion processed primarily along the dorsal pathway and hue along the ventral pathway (24, 25). Functionally these features have also often been referred to as emblematic of the distinction between "what" and "where" processes, or of vision for perception vs. action (26). Yet, despite these physiological and functional differences, both color and motion coordinates can be described by three canonical cardinal axes (Fig. 1*A*). For motion, these axes represent left–right, up–down, and near–far, while for color, these axes in classical color-opponent theory represent red–green, blue–yellow, and white–black. To determine whether this shared dimensionality reflects common underlying perceptual representations, we compared motion judgments in the two-dimensional frontoparallel plane (at constant depth) with hue judgments in the two-dimensional chromatic plane (at constant luminance) using functionally equivalent psychophysical tasks. Measurements from these tasks show that for both stimulus attributes, there are large and systematic differences in how observers judge directions in the plane, and analyses of these differences point to fundamentally different principles for their perceptual representation.

Results

Perceptual Judgments of Hue and Motion Direction. To measure judgments of hue, 46 observers completed a hue-scaling task (Experiment 1) in which they described each stimulus by decomposing the perceived hue into percentages of each primary hue (red, green, blue, and yellow) for a total of 100% (e.g., describing an orange hue as a proportion of red and yellow) (Fig. 1*B*). Since this task was developed, it has become a standard method for measuring hue percepts (27). On each trial the chromatic angle of the stimulus was chosen pseudorandomly from 36 possible stimulus directions created by sampling the L vs. M and S vs. LM cone-opponent plane at 10° intervals (Fig. 1*C*). Again, these axes correspond to the cardinal directions underlying color coding at early postreceptoral stages. The stimulus was presented as a pulsing (0.5 s on/1 s off) uniform square subtending 2° of visual angle that continued until the observer entered their response. Four trials were presented for each stimulus angle.

To measure judgments of motion direction in an analogous way, we designed a motion-scaling task (Experiment 2). An independent set of 48 observers participated in motion scaling in which they described the motion direction of an array of moving dots using percentages of the four primaries (up, down, left, and right) for a total of 100% (Fig. 1*B*). On each trial, the motion stimulus was again chosen pseudorandomly from the 36 corresponding directions (Fig. 1*C*) and presented as a pulsing (0.5 s on/1 s off) circular array of moving dots within a 2° circular visual field, that again continued to pulse until a response was made. Four trials were presented for each stimulus angle.

To calculate the *scaling function* for each observer, the percentages assigned to each stimulus angle were converted into a perceived angle within a 2D space with cardinal axes defined by the primaries of each task (Fig. 1*D*). For example, if the stimulus was scaled as appearing 50% red and 50% blue, then this corresponded to a perceived angle of 45° in the red–green (0 to 180°) and blue–yellow (90 to 270°) perceptual plane. Thus, the scaling function represented the perceived angle as a function of the stimulus angle, and for each observer was calculated based on the mean of their four responses (Fig. 2*A*).

To verify that for each task that the differences between observers were significant (22), we compared the magnitude of the inter- and intraobserver variability, after arcsine-transforming the scaling functions to normalize the variance. The within-observer variability was similar for both tasks, with an average SD of roughly 5°. One-sided Mann–Whitney *U* tests confirmed that SDs across mean settings between observers were larger than SDs across trials within observers for both the hue- [$Z(46) = -7.27$, $P < 0.001$] and motion- [$Z(48) = -4.99$, $P < 0.001$] scaling tasks (Fig. 2*B*). On average, the variability between subjects was 2.8 (hue) or 2.1 (motion) times greater than the within-subject variation. We further confirmed significant test–retest reliability [Spearman's rank-order correlation: $\rho(1,654) = 0.52$, $P < 0.001$; $\rho(1,726) = 0.42$, $P < 0.001$], and split-half reliability [$\rho(1,654) = 0.63$, $P < 0.001$; $\rho(1,726) = 0.99$, $P < 0.001$] for hue and motion, respectively. These analyses suggest that the individual differences in the data are not due to only noise across trials or measurements in the scaling tasks, and might instead reflect systematic differences in how observers rate the perception of these stimulus attributes.

Simple comparisons of the hue- and motion-scaling data show distinct patterns of responses between the tasks. For example, the SDs (calculated on arcsine-transformed percentages) as a function of stimulus angle showed stronger evidence for privileged or unique directions in motion scaling, revealed by stimulus angles where there is little variability across observers (Fig. 2*B*). For motion,

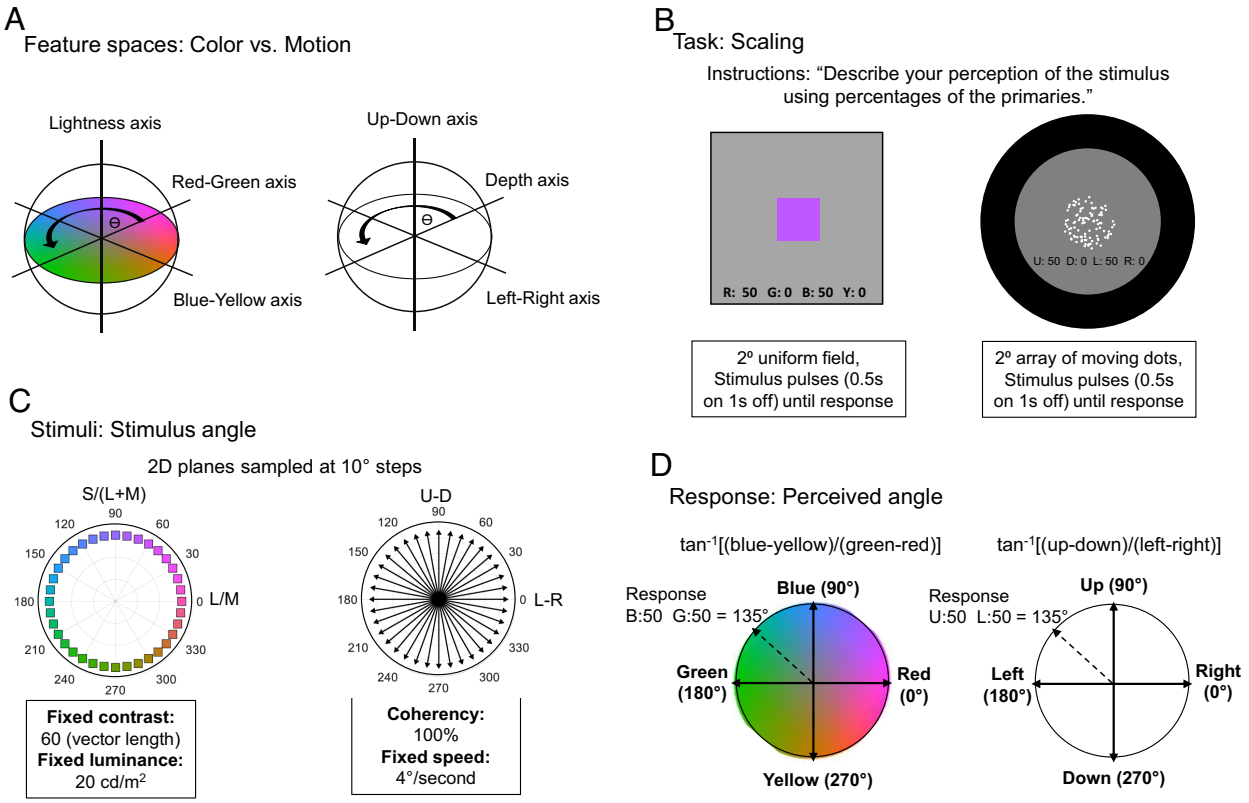

Fig. 1. Experimental design. (*A*) The feature spaces for color (*Left*) and motion (*Right*) share a dimensional structure in which three fundamental, opponent, perceptual axes are sufficient to fully define the stimuli. (*B*) To reveal whether color and motion percepts are also coded by common principles, we measured observers' judgments of color and motion using analogous scaling tasks. For this task, observers are instructed to describe each stimulus as a combination of the primaries (color: *Red*, *Green*, *Blue*, and *Yellow*; motion: *Right*, *Up*, *Left*, and *Down*). They assigned percentages to the primaries to report the combination that aligned with their percept. For hue scaling, on each trial the observer was presented with a 2° uniform square of color that pulsed continuously (0.5 s on/1 s off) until their response was made. For motion scaling, on each trial the observer was presented with a 2° circular array of randomly distributed moving dots that again pulsed continuously (0.5 s on/1 s off) until their response was made. (*C*) The stimulus presented for each trial was chosen randomly from 36 possible directions for either color or motion. For color, these 36 directions were defined by sampling the circumference of a circle of fixed contrast and luminance in the LvsM and SvsLM chromatic plane at 10° steps. For motion, the 36 directions were defined by sampling the circumference of a circle of fixed velocity and depth at 10° steps in the frontoparallel plane. The array of dots moved 100% coherently at a fixed speed in the chosen direction on each trial. Observers scaled each of the 36 angles four times. Independent groups of observers completed the hue and motion scaling tasks. (*D*) The response for each stimulus angle was then converted to a perceived angle for subsequent analyses. A perceived angle corresponds to the angle within the perceptual space used to scale the stimuli (i.e., *Red-Green* and *Blue-Yellow* axes, or *Left–Right* and *Up–Down* axes). Response of 50% *Blue* and 50% *Green*, or 50% *Up* and 50% *Left*, would therefore correspond to a perceived angle of 135°. A scaling function was calculated for each observer, which defined the mean perceived angle at each stimulus angle for each observer.

these privileged directions aligned with the cardinal axes. For hue, these privileged directions corresponded with typical red, blue, green, and yellow loci. However, the SDs for these primary hues are higher than those for the primary axes for motion, and did not always reflect the minima in the observer differences (e.g., there is greater inter-observer agreement for purple than red). These differences therefore provide stronger evidence for privileged directions in motion percepts than hue percepts.

A more dramatic difference is evident in the correlation matrices for the two datasets (Fig. 2*C*). For hue, the responses to different stimulus angles primarily showed only positive correlations and only for nearby angles, along the diagonal of the matrix. In contrast, for motion there were consistent positive and negative correlations among both near and widely separated angles. The robust correlation structures for both datasets indicated fewer, latent variables influencing observers' responses. In the following analyses, we used factor analysis to reveal these latent sources of variability for each task (23, 28, 29). Factor analysis is a standard statistical procedure for extracting the patterns of variation underlying an observed set of measurements. The basic premise is that correlated observations reflect the influence of one or more common factors, while uncorrelated observations depend on different factors. From this analysis, we could therefore compare the number and nature of the processes underlying hue and motion direction judgments and what these might imply about the perceptual representations of these attributes.

Many Factors Spanning Different Local Regions of the Hue Circle Underlie Individual Differences in Hue Judgments. The factor analysis of hue-scaling functions revealed nine systematic factors (defined empirically as factors with significant loadings on two or more adjacent items, suggesting they represent meaningful variation and not noise, see *Methods: Systematic Loadings* for more information). Each of these factors had high loadings for a narrow range of stimulus angles and with loadings of the same sign (Fig. 3*A*). The loadings for these factors across stimulus angle thus showed a single peak and could be reasonably approximated by a simple Gaussian (*SI Appendix*, Table S2). Other common methods for determining the factor solution suggested similar results (*SI Appendix*, Table S1 and Fig. S1). Together, the nine factors accounted for 76% of the total variance and much of the correlation structure, as evidenced by the lack of variance magnitude and structure in the residual matrix (*SI Appendix*, Fig. S2). These factors are similar in number and nature to those revealed in previous analyses of hue-scaling functions based on a smaller sample of observers (17, 21) suggesting that this factor pattern is robust.

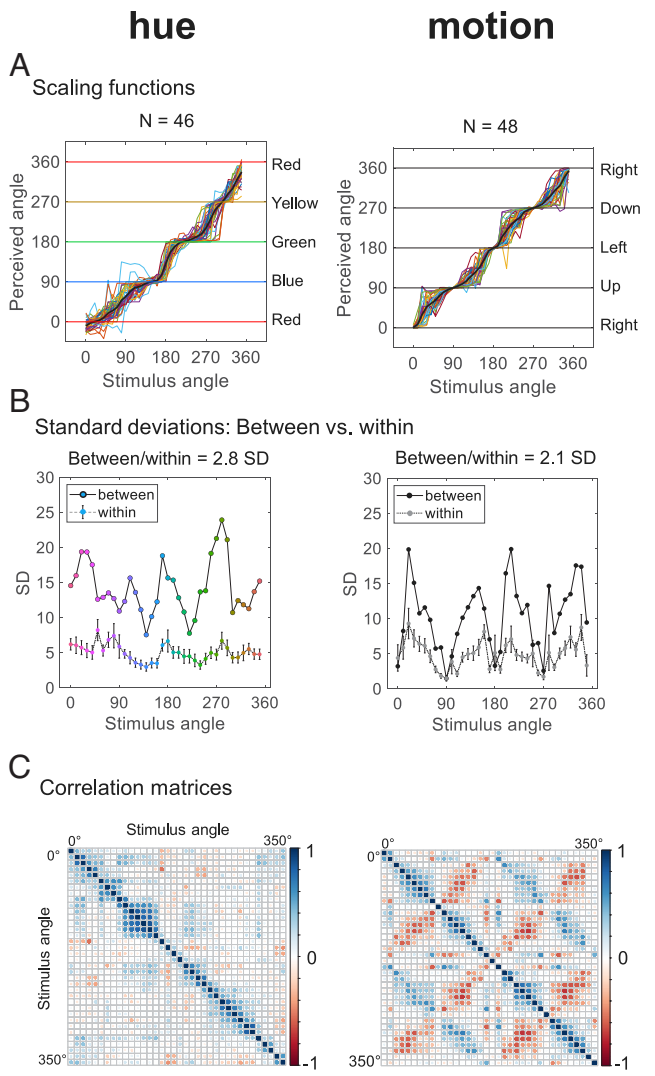

Fig. 2. Comparisons of the pattern of individual differences in hue- (*Left*) and motion- (*Right*) scaling functions. (*A*) Scaling functions are plotted for 46 and 48 observers for the hue- and motion- scaling tasks, respectively (colored lines show the scaling function for each observer). The mean scaling function for each task is indicated in black. Horizontal lines indicate the primaries for each task, which correspond to the cardinal directions of the perceived angles. (*B*) In order to validate the individual differences, we compared the magnitude of the SD in settings *within* an observer across trials and *between* observers for each stimulus angle. The SDs were computed on the arcsine-transformed scaling data to correct for variance compression at the extremes. Solid lines indicate the SD between observers and the dotted line the average SD within observers across trials. Error bars represent +/-1 SEM, and could only be calculated for the values within observers. Ratios above the plots correspond to the ratio of the SD between observers to the average SD within. These values suggest that the individual differences in these data cannot be attributed to measurement variability across trials. (*C*) The correlation structure in the responses between tasks.

An independent components analysis of the hue-scaling functions also revealed similar processes (*SI Appendix*, Fig. S3). The large number of narrow factors is inconsistent with variability in a small number of canonical primary hues, while the unimodal loadings are inconsistent with an opponent representation (which would predict a bipolar pattern of loadings). Thus, in contrast to the two chromatic dimensions assumed by classic color-opponent theory, these results point to a representation with limited evidence for privileged axes and no direct evidence for opponent axes.

To visualize the relation between this observed pattern and the pattern predicted from different models of color appearance, we calculated the predicted correlation matrices for scaling functions based on simulated sources of individual variability (Fig. 4). These sources of variability included two broadly varying bipolar axes (i.e., a conventional opponent code), four broadly varying unipolar processes (i.e., a rectified opponent code so that "red", "green", "blue", and "yellow" could vary independently), and eight more narrowband or localized processes (so that each process spanned a narrow range of chromatic angles). For each case, we randomly varied the peak angle and sensitivity of each process and then constructed the corresponding hue-scaling functions from their combined responses. Comparisons of these matrices to the observed correlations in Fig. 2*C* show that for hue the pattern of variability is qualitatively similar to the model with eight localized sources of variability, while different from the opponent or rectified-opponent models. In particular, the opponent model predicts both positive and negative correlations (since for example, "red" and "green" are negatively correlated), which are not evident in the observed pattern. The rectified-opponent model also fails because the "red" or "green" processes, etc., are too broad in their bandwidth to account for the more localized correlations in the observed matrix. Thus, even though observers were required to decompose the hue in terms of "red–green" and "blue–yellow" responses, these primaries do not emerge as the source of individual differences in the hue judgments. Adding additional sources of variability to the models (e.g., nonlinearities in the responses or categorical biases) similarly fail to account for the narrow, localized factors that emerge for hue scaling (17). Instead, the observed factors show that different narrow hue regions vary largely independently, suggesting that they are not tied to a common underlying variation in a set of perceptual primaries. Finally, note that it is the narrow range of each process, and not the number of processes, that predicts the narrow range of the observed factors. That is, replacing the eight discrete processes with a continuum (e.g., many processes with peaks spaced evenly along the spectrum and equal bandwidths) would lead to less discrete clustering in the correlations, but would retain the narrow, localized pattern of the correlations.

We next examined the possible basis for the hue-scaling factors. One possibility we considered is that the eight factors might reflect variations in how observers represent hues in terms of eight hue categories (the four pure hues (e.g., red) and the four binary hues (e.g., orange)). In this case, the factors could capture differences in the focal stimuli (e.g., 100% red or 100% orange), or in the boundaries between them (e.g., 50% red and 50% orange). We estimated these foci and boundaries directly from the observed hue-scaling functions. For example, a perceived angle of 0° corresponds with pure red, while a perceived angle of 45° corresponds with the boundary between red and blue. We then examined how observers' hue categories were related to their scores on each factor. Factor scores indicate the contribution of the latent factors to each observer's hue-scaling responses, and therefore determining the variables with which these scores are correlated can help identify the underlying processes (29). Using nearest-neighbor interpolation, we directly estimated the stimulus angles corresponding to 16 different perceptual angles in the observers' hue-scaling functions, sampling the perceptual space in steps of 22.5° to estimate the loci and boundaries between primary (red, green, blue, and yellow) and binary hues (orange, purple, blue–green, and yellow–green). We estimated factor scores using Thurstone's (29) least squares regression approach, which is recommended for use with PC factor analysis methods (28). Spearman's rank-order correlations between observers' settings for the loci and boundaries and the factor scores suggest a stronger relationship between the factors and the category boundaries than the category exemplars (Fig. 3*B*). Specifically, except for factor 1 and red, none of the other factors align with the foci

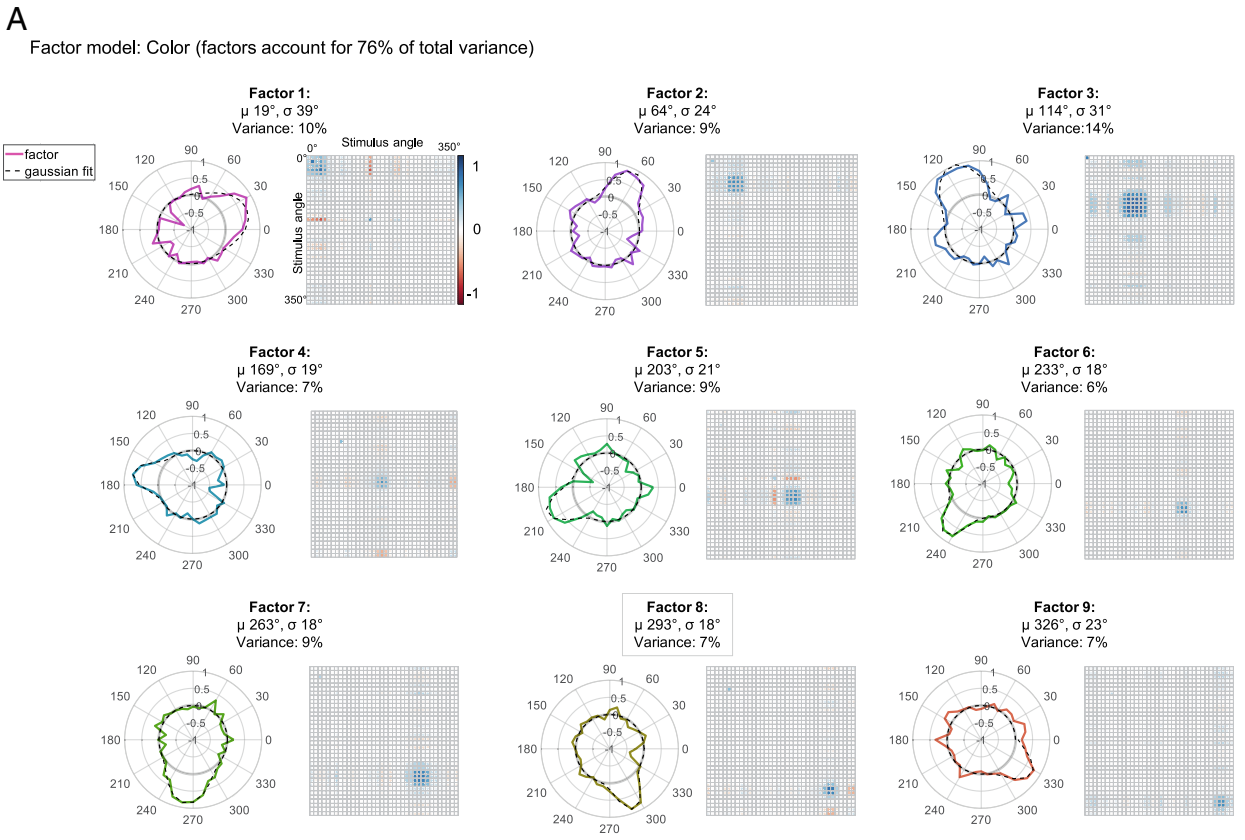

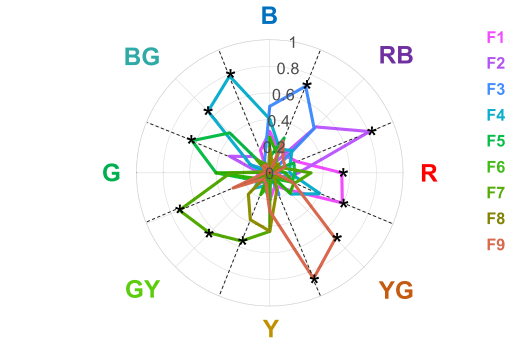

Fig. 3. A factor analysis of hue-scaling functions revealed many processes that accounted for the variance for narrow regions of color space. (*A*) The factor model for color consisted of nine systematic (see *Methods*) factors, which together accounted for 76% of the total variance. Panels show the loadings for individual factors and the correlation structure of the original data accounted for by each factor (estimated by $L_f L_f' + U_f$, where L is the matrix of factor loadings and U is the matrix of uniqueness). Each factor loading plot shows the value of the factor loading (radius), which represents the proportion of variance in the responses to each stimulus accounted for by each factor (a light gray line indicates a radius or loading of zero in each plot for reference of the neutral point). The factors are colored to indicate the hue category corresponding to the maximum loading. The black dotted line in each factor plot corresponds to the best-fitting Gaussian for that factor, with the peak and bandwidth as reported above each plot. The total variance accounted for is also reported for each factor. (*B*) The scores for each factor were then analyzed to identify the process defined by each factor. Category loci and boundaries inferred by nearest-neighbor interpolation were correlated with the factor scores in order to determine their relationship. The polar plot shows that the observers' category boundaries are more strongly correlated with the factors than the loci themselves (starred comparisons indicate significant correlations ($P < 0.05$, after correction for multiple comparisons)).

for the primary or binary hues. Rather they are generally associated with the boundaries between primary and binary hues (Fig. 3*B*). These analyses thus suggest that the individual differences in hue scaling are primarily around the angles where observers demarcate the boundaries between different color categories rather than at the category exemplars themselves. This finding aligns with research showing less agreement in the color terms assigned to category boundaries than category centers (30, 31). That the highest correlations between the factor scores and scaling responses lie at the boundaries between primary and binary categories further suggests that a representation in terms of underlying principal hue axes is not sufficient to account for individual differences in hue-scaling judgments.

In addition to exploring the relation of the factors to different color categories, we tested for evidence of "categorical coding" by directly assessing the hue-scaling functions. This coding predicts that the reported hues of two stimuli that both fall within a category should be more similar than between two stimuli that instead fall in different categories and predicts that the hue-scaling functions should exhibit discrete, step-like changes between categories rather than smooth variation (32). We estimated the degree of categorical bias exhibited by each observer by fitting a weighted average of a linear and step function to their scaling function (32). However, this revealed relatively weak categorical biases in the settings, with the degree of the bias largely uncorrelated across the different hue categories (*SI Appendix*, Fig. S9). Thus, although the factors tended to align with color category boundaries, percepts varied more continuously than they did categorically, suggesting only a slight influence of categorical coding.

A Small Number of Broad, "Opponent-Like" Factors Underlie Individual Differences in Motion Judgments. We applied the same set of analyses to the motion-scaling functions to determine the underlying representational structure for motion percepts. The

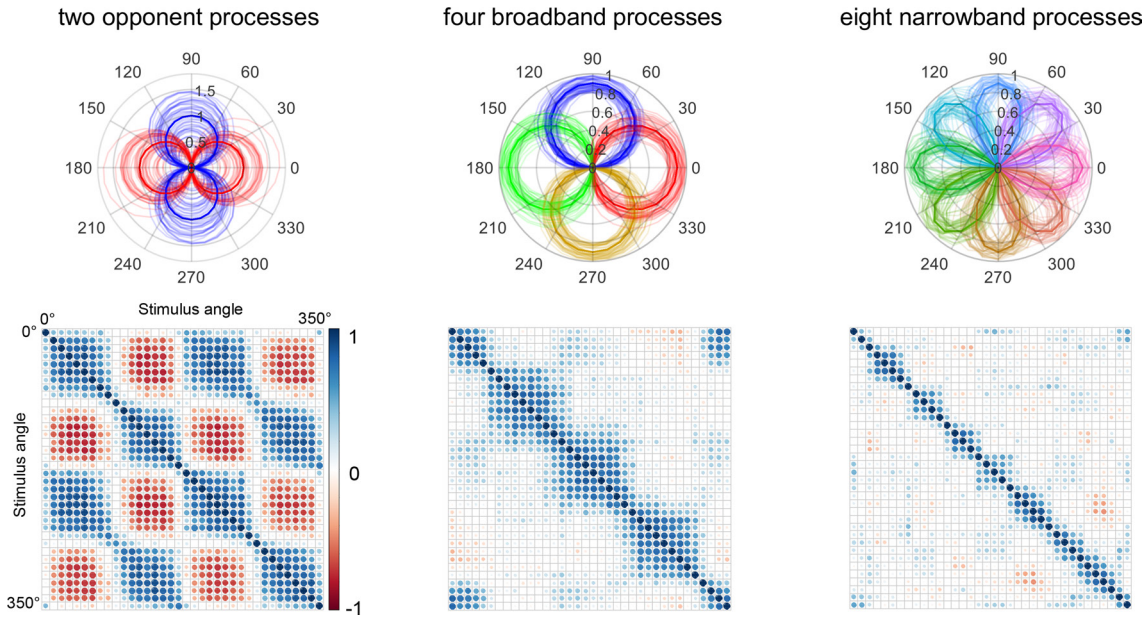

Fig. 4. Hypothetical models and correlation matrices for simulated sources of individual variability in hue scaling. *Upper* plots show hypothetical opponent (*Left*), broadband (*Middle*), or narrowband (*Right*) processes. Each process was independently varied in peak angle and magnitude. The corresponding hue-scaling functions were derived by combining the responses across processes perturbed with Gaussian noise similar in magnitude to the observed within-observer variability. Lower plots show the resulting correlations across stimulus angles calculated from the simulated variations in hue-scaling functions predicted by each model.

factor analysis of motion-scaling functions revealed three systematic factors that were each periodic and therefore qualitatively different from the unimodal factors observed for hue (Fig. 5*A*). Other common methods for determining an adequate number of factors suggested similar results (*SI Appendix*, Table S1 and Fig. S4). Together, the three factors accounted for 52% of the total variance and much of the important correlation structure, as evidenced by the lack of variance magnitude and structure in the residual matrix (*SI Appendix*, Fig. S5). An independent components analysis of the motion-scaling functions also revealed a similar component pattern (*SI Appendix*, Fig. S6). Given the lower amount of variance accounted for by systematic factors for the motion compared with the hue factor model, we also explored a nine-factor model for the motion-scaling data to understand whether the additional factors were meaningful and similar to those for hue. The nine-factor model for motion accounted for 76% of the total variance and revealed that the additional factors were extremely narrow, with high loadings near the cardinal axes (*SI Appendix*, Fig. S7). [They, therefore, did not meet the "systematic" criterion of consistent loadings across adjacent angles for inclusion in the initial analysis, but could nevertheless reflect meaningful sources of interobserver variation (i.e., not just measurement noise).] The overall factor pattern for motion is thus very distinct from that for hue. In particular, unlike the multiple localized factors mediating differences in hue judgments, the pattern of individual differences in motion percepts was consistent with a code in terms of global, cyclical processes. This correspondence for motion can be seen by the qualitative similarities between the correlation matrix predicted by the opponent-process model (Fig. 4) and the observed correlation matrix for the motion (Fig. 2*C*).

We again examined the processes to which the motion factors might correspond. In this case, we approximated the factor loadings by fitting a sawtooth function to characterize the phase and frequency for each factor (Fig. 5*A*). (Because of asymmetries in the loadings, a sawtooth provided a better fit than a sinewave variation; see *SI Appendix*, Table S3). The phase and frequency of the best-fitting sawtooth wave for each factor is shown in Fig. 5*A*. Based on these values, the factors corresponded to variations in judgments of motion directions that were intermediate to the vertical and horizontal axes. The first two factors had a frequency of one cycle (per 360°) and peak loadings near 60° to 240° (Factor 1), or 30° to 210° (Factor 2). The third factor instead had a frequency of two cycles, again peaking at intermediate angles.

One possibility we considered was that all three factors in fact reflected the special status of the horizontal and vertical axes but depended on individual variations in the relative salience or sensitivity to these cardinal directions. Note this predicts that the factors should have the strongest loadings for off-axis directions (reflecting individual differences in the relative weights given to the vertical vs. horizontal components), while weakest loadings for the cardinal axes (where the motion depends on only the horizontal or vertical component). This pattern is consistent with the markedly lower interobserver variations in motion scaling for the horizontal and vertical compared with intermediate axes (Fig. 2*B*), and with the finding noted above that the variations along the cardinal axes showed up as additional, narrowly-tuned factors accounting for less variance (*SI Appendix*, Fig. S7). Thus, while such factors are not aligned with the cardinal axes, they would nevertheless reflect variations that are tied to these axes. To formally assess this, we computed the relative amplitude between the horizontal and vertical components of each observer's motion-scaling function (*SI Appendix*, Fig. S8). Horizontal/vertical ratios across participants revealed a general trend consistent with decreased scaling for the horizontal axis relative to the vertical axis [as documented in many previous studies (33)], with considerable individual differences in the magnitude of this bias (mean = 2.22, SD = 1.67) (*SI Appendix*, Fig. S8). To examine whether this asymmetry was the basis for the first two factors, we computed a knock-out of this horizontal/vertical bias by equating the amplitude of the horizontal and vertical components of each observer's scaling function. We then factor-analyzed the scaling functions after this correction, revealing a solution where the first two factors had now dropped out (*SI Appendix*, Fig. S8). These results suggest that individual differences in the relative weights that observers assign

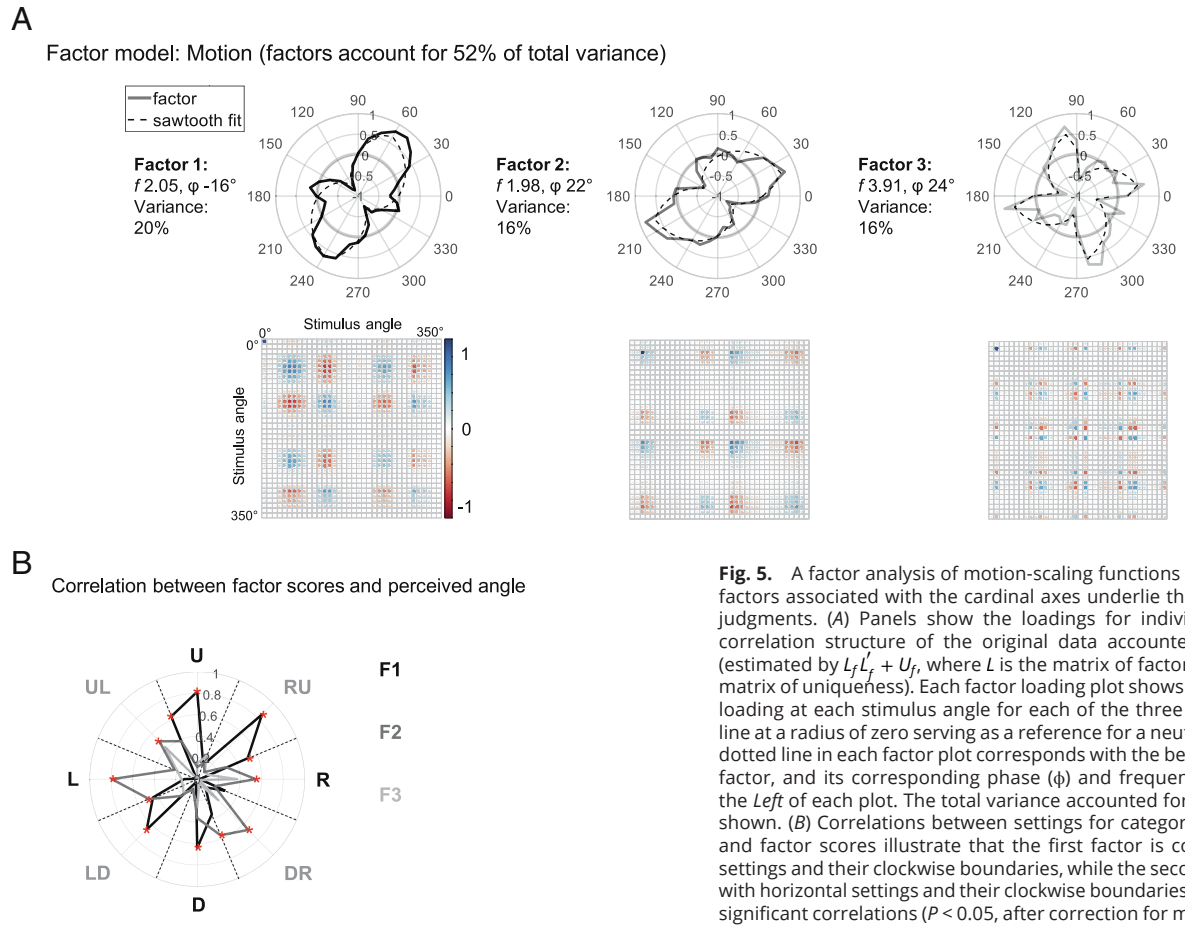

Fig. 5. A factor analysis of motion-scaling functions revealed that periodic factors associated with the cardinal axes underlie the variability in motion judgments. (*A*) Panels show the loadings for individual factors and the correlation structure of the original data accounted for by each factor (estimated by $L_f L_f' + U_f$, where L is the matrix of factor loadings and U is the matrix of uniqueness). Each factor loading plot shows the value of the factor loading at each stimulus angle for each of the three factors (with the gray line at a radius of zero serving as a reference for a neutral loading). The black dotted line in each factor plot corresponds with the best Sawtooth fit for that factor, and its corresponding phase (ϕ) and frequency (*f*) are reported to the *Left* of each plot. The total variance accounted for by each factor is also shown. (*B*) Correlations between settings for category loci and boundaries and factor scores illustrate that the first factor is correlated with vertical settings and their clockwise boundaries, while the second factor is correlated with horizontal settings and their clockwise boundaries [red asterisks indicate significant correlations ($P < 0.05$, after correction for multiple comparisons)].

to the horizontal and vertical cardinal axes account for much of the variability in motion direction judgments (*SI Appendix*, Fig. S8). (Factors with a similar pattern of loadings were not observed for hue, and when a similar knock-out analysis was applied to the hue-scaling functions, it did not lead to a clearly interpretable result; *SI Appendix*, Fig. S8). The third factor for motion (which again had twice the frequency of variation as the first two factors) survived the knockout. However, this factor could also reflect variations relative to horizontal and vertical, if it corresponded to variations in the degree of bias toward the cardinal axes (e.g., in the tendency to overweight the contribution of the nearest cardinal direction to the perceived motion). However, additional modeling would be required to verify whether this bias aligned with this factor.

Finally, as with the hue analyses, we also investigated the extent to which the motion factor scores corresponded with category boundaries or foci. Again, using nearest-neighbor interpolation, we directly estimated the stimulus foci from the mean motion-scaling function of each observer for the same sixteen perceptual angles, sampling the underlying perceptual space in steps of 22.5°. We then estimated the factor scores using the same procedure reported above. Categorical biases in this case were again weak (and significantly weaker than for hue; *SI Appendix*, Fig. S9). In contrast to hue, Spearman's rank-order correlations between observers' settings for the loci and boundaries and the factor scores suggested a stronger relationship between the motion factors and settings at the primaries and their boundaries (Fig. 5*C*). Specifically, settings for the vertical axis and the boundaries nearest to and clockwise from it correlate significantly with the scores for the first factor, and settings for the horizontal axis and boundaries nearest to and clockwise from it correlate significantly with the scores for the second factor. These analyses further suggest that the motion factor pattern depends on latent processes related to the cardinal axes of visual space, further distinguishing the pattern of responses for motion and hue.

Discussion

Our analyses reveal that, while the coordinates for all colors or motions can be described in three-dimensional spaces, the perceptual representations of the attributes of hue and motion direction are fundamentally distinct. Moreover, these differences emerged using similar perceptual judgments and analyses. This is important, because it supports the conclusion that the patterns of individual differences revealed for hue or motion direction percepts are not simply a property of the experimental task or analyses used to assess them, but instead reflect actual differences in their visual representations.

In this study, we did not extend the comparison of color and motion coding to their full three-dimensional gamut. There is some evidence to suggest broad tuning along the third dimension of color, brightness (8, 10, 34), with the potential for new color categories to emerge as this dimension is modulated (e.g., between luminance increments and decrements) (35); and though the coding principles in the context of the third dimension of visual motion (depth), are quite complicated, there is some suggested evidence for a representation that references spatial geometry, especially in the context of actions such as accurate object tracking and interaction (36, 37). However, this evidence does not come directly from the psychophysical scaling tasks used here, and we suggest that extending the reported analyses to all three dimensions is an

important future direction for a comprehensive comparison of the perceptual representations underlying these visual features.

For the two-dimensional coordinate frames we examined, the response patterns for hue reflected multiple, relatively local processes with no direct evidence for opponent axes or a small set of (four) perceptual primaries, corroborating previous hue-scaling studies (17, 21). Moreover, there was little evidence for a superordinate status for certain hues. The interobserver variance in the settings did tend to be lower for the primaries. However, this status itself may not reflect a special representation for red–green and blue–yellow since these hues were the constructs used to scale the stimuli, and previous research suggests that this difference in variability is not present under a different task (18). It would be instructive in future studies to determine if the pattern of variation would change if a different set of primaries, for example corresponding to the binary hues (19), were used to rate the stimuli.

For motion direction, the response patterns instead depended on three latent processes with broader and more global influences on the settings. These primarily captured the large interobserver variations at angles intermediate to the cardinal directions but were consistent with variations in the weights given to the cardinal axes, and thus to a representation associated with them (Fig. 4). Moreover, the factor scores for motion were more clearly associated with the cardinal axes and the boundaries between them, again pointing to an underlying representation based on a privileged pair of axes. Thus, judgments about the direction of stimuli within the two "spaces" appeared to draw on very different representations.

Where do these differences in representation emerge? One possibility is that they reflect differences in how hue and motion are encoded by the visual system. However, this leads to the paradox that the variability in scaling does not reveal opponent processing for hue, while for motion it does. Moreover, this interpretation is inconsistent with the functional similarities between the early cortical representation of color and motion. Like many other visual attributes, at this level cells show preferences for a wide variety of chromatic angles (38) or motion directions (39), suggesting a common representational format in terms of a multiple-channel or population code (40). Given these analogous encoding schemes, it is reasonable to assume that the factors found are not directly associated with the tuning or number of encoding mechanisms. Instead, the different factor patterns for hue and motion may be more closely related to how information from the population responses is decoded. The multiple, narrow factors for hue might suggest that the decoding is without obvious reference to an underlying coordinate frame, potentially so that hues are represented more like qualitatively different objects than quantitatively different directions. This might be related to the finding that it can be difficult to intuit the coordinate relationships between different colors (e.g., which colors are complementary or orthogonal (15)), and to the way we communicate about color using several basic color terms (41). For motion, the factor pattern is instead consistent with a decoding that seemingly preserves information in terms of an underlying reference frame. This might be reflected in that, for motion, it seems more intuitive to compute the coordinate relationships between different trajectories or points in space, that perceptual errors are often biased toward the cardinal axes (42, 43), and that, unlike color, we often communicate about different motion directions with terms referenced to the cardinal axes (e.g., upper left, or north-northwest).

Why are the decoded representations for hue and motion direction different? As noted in the Introduction, the basis for the common perceptual dimensionality of these attributes is distinct. For motion it may be associated with the geometry of physical space (or alternatively, with prominent features of visual space such as the vertical axis of gravity and a horizontal ground plane). In contrast, for hue it derives from comparisons of the three samples for spectral stimuli provided by the cones (with both hue and motion direction also influenced by the stimulus context). Thus, unlike perceptual color space, perceptual motion space has external validity, at least as far as how the observer interacts with it. For example, for spatial navigation and action, it is arguably critical to sense and interpret the spatial information meditated by vision, for example, to reach out to the location one sees. In contrast, for attributes like hue, there is not a corresponding ground truth in what is red, nor is there evidence that one is necessary as long as the percepts convey consistent information. While we have focused only on hue and motion direction, it would be valuable in future work to assess whether these differences reflect more general strategies for how information is used within visual processing streams for carrying inherently spatial vs. nonspatial information. In any case, our results suggest that conventional models of color appearance, based on an underlying pair of primary hue dimensions, may provide a poor account of the actual perceptual representation of hue.

Methods

Participants. The observers included 46 (30 females and 16 males; mean age ± SD = 22.87 ± 5.01 y) and 48 (32 females and 16 males; mean age ± SD = 24.5 ± 6.75 y) undergraduate and graduate students at the University of Nevada, Reno (UNR), for the color and motion experiments, respectively (with independent groups of observers tested for the two experiments). Data from 26 of the participants in Experiment 1 are from Emery et al. (17). Data from two and four additional observers were excluded for the color and motion experiments, respectively, due to inconsistencies in their settings across trials (see *Methods: Data Preprocessing*). All procedures followed protocols approved by UNR's Institutional Review Board, and participation was with written informed consent, with students offered course extra credit for their participation. The observers were screened for self-reported normal or corrected-to-normal vision, and for the color experiments were screened for normal color vision using the Cambridge Colour Test (Cambridge Research Systems).

Color Apparatus and Stimuli. Stimuli were presented on a SONY Multiscan Trinitron 500PS CRT monitor controlled by a Cambridge Research System ViSaGe MKII Stimulus Generator, providing 12-bit resolution per gun. The monitor was calibrated with a Photo Research 655 spectroradiometer and gun outputs were linearized via gamma correction. The stimuli and background were presented at a constant luminance of 20 cd/m^2. The gray (Illuminant C, CIE x: 0.31, y: 0.316) background filled the (41.5 × 31.5 cm) monitor screen, subtending 11.7° × 9° of visual angle. The test chromaticities were based on a variant of the MacLeod–Boynton (44) chromaticity diagram, scaled to equate sensitivity along the L vs. M and S vs. LM axes based on chromatic thresholds (45):

$$L\ vs.\ M = \left(l_{mb} - 0.6568\right) * 2{,}754,$$

$$S\ vs.\ LM = \left(s_{mb} - 0.01825\right) * 4{,}099.$$

L vs. M and *S vs. LM* are the cardinal axes of the underlying chromatic, cone-contrast space. The variables l_{mb} and s_{mb} are the *l* and *s* (or *r* and *b*) coordinates in MacLeod–Boynton space, and their contrasts are relative to the chosen achromatic point in the space (0.6568, 0.01825, the MacLeod–Boynton coordinates for Illuminant C), which are then scaled to approximate equal multiples of threshold. (Note that compared with motion, this scaling for color is somewhat arbitrary, and cannot be equated directly to the strength of the motion stimuli. Changing the scaling would change the relationship between hue and stimulus angle but would not change the general pattern of the correlations in hue percepts across stimulus angles. There are also a number of other ways that we could not formally equate the motion and color stimuli, e.g., one is dynamic while the other static.)

Stimuli within this space were defined by a vector with a direction corresponding to the chromatic angle and a length corresponding to chromatic contrast. Thirty-six test chromaticities were defined by sampling the space at 10° steps from

10 to 360° at a fixed contrast of 60 (Fig. 1*C*). Each chromaticity was presented foveally as a 2° uniform square pulsing for 0.5 s, with a 1s return to the gray background between each pulse. For each trial, the same stimulus was repeated until the observer made their setting. Observers viewed the display binocularly in an otherwise dark room from 200 cm.

Hue-Scaling Task and Procedure. In the hue-scaling task, observers were instructed to report their perception of the presented hue as a percentage of red, green, blue, or yellow in increments of 5% (17). Observers were instructed that the sum of their assigned percentages must equal 100% with the constraint that red and green could not be used simultaneously to describe a hue, nor blue and yellow (46). These hue terms are rarely paired together when observers are free to do so (47), reflecting the common notion that the opponent pairs represent mutually exclusive sensations.

Perceived angles were calculated for each response by converting the percentages to an angle in a red–green (0 to 180°) and blue–yellow (90 to 270°) perceptual space (Fig. 1) by the following formula (16, 17):

$$Perceived\ angle = \tan^{-1}\left(\frac{(blue - yellow)}{(red - green)}\right).$$

For example, if an observer responded that a stimulus appeared 50% blue and 50% green, this would correspond to an angle of 135° in this space, or an equally balanced blue–green hue. The perceptual angles rather than the percentages were used for subsequent analyses.

Observers participated in two sessions of hue scaling. Within each session, the 36 stimulus angles were presented two times in a pseudo-random order for a total of four trials per stimulus angle. Each session lasted less than one hour, and the two sessions were completed on separate days within a week. On the first day and before starting the experiment, observers completed a set of practice trials with the experimenter to ensure they understood the task. There were twelve practice trials which consisted of scaling six stimulus angles twice (sampling the color space in 60° steps from 12 to 252° to ensure no colors in the practice set were included in the experiment). Following the practice trials, observers adapted to the monitor background in the otherwise dark room for 1 min before starting the first trial of the experiment. On each trial, the stimulus continued to pulse until the observer entered their response, at which point the next stimulus angle was presented. Between each of the two blocks within each session (each block consisted of the full set of 36 stimulus angles), the observer was presented with a gray intermission screen and pressed a key when they were ready to begin the next block.

Motion Stimuli and Apparatus. Stimuli were presented on an identical calibrated monitor, with a black, circular frame placed in front to remove horizontal and vertical edge information from the screen (Fig. 1*B*). A gray background with a luminance of 20 cd/m^2 was presented on the monitor and viewed through the screen aperture which subtended 9° of visual angle. Each stimulus was a circular array of moving dots subtending 2° of visual angle presented foveally and in the center of the background (Fig. 1*B*). The array consisted of 100 white dots, with each subtending 0.1° of visual angle, and distributed randomly within the circular aperture at the initialization of each trial. On each trial, the dots moved 100% coherently in one of 36 test directions corresponding to angles at 10° steps in the frontoparallel plane. Dot density remained constant during the presentation, consistent with viewing a uniform array through a 2° aperture. (Fig. 1). The dots moved along the chosen direction at a fixed speed of 4° per second for 0.5 s, with a 1 s return to the gray background between each presentation. During a trial, the same stimulus was repeated until the observer completed their setting. Observers viewed the display binocularly in an otherwise dark room from a distance of 200 cm.

Motion-Scaling Task and Procedure. For the motion-scaling task, observers were instructed to report the perceived motion direction as a percentage of up, down, left, or right. Responses could be varied in increments of 5%, and observers were instructed that the sum of their assigned percentages must equal 100%. We again added the limitation that right and left are mutually exclusive, as are up and down. Perceived angles were calculated for each response using a procedure analogous to hue scaling, by converting the percentages to an angle in a right–left (0 to 180°) and up–down (90 to 270°) perceptual space (Fig. 1) by the following formula:

$$Perceived\ angle = \tan^{-1}\left(\frac{(up - down)}{(right - left)}\right).$$

Therefore, if an observer responded that a hue appeared 50% up and 50% left, this would correspond to an angle of 135° in this space. The perceived angles rather than the percentages were used for subsequent analyses.

As in hue scaling, observers participated in two sessions of motion scaling. Within each session, the 36 stimulus angles were presented two times in a pseudo-random order for a total of four trials per stimulus angle. Each session lasted less than one hour and the two sessions were completed on separate days within a week. On the first day and before starting the experiment, observers completed a set of practice trials with the experimenter to ensure they understood the task. There were 12 practice trials which consisted of scaling six angles twice (sampling the motion space in 60° steps from 12 to 252° to ensure no motion directions in the practice were included in the experiment). Following the practice trials, observers adapted to the background in an otherwise dark room for one minute before starting the first trial of the experiment session. On each trial, the stimulus continued to pulse until the observer responded, at which point the next randomly chosen direction was presented. Between each of the two blocks within each session (each block consisted of the full set of 36 stimulus angles), the observer was presented with a gray intermission screen and pressed a button to initiate the second block of trials when they were ready.

Data Preprocessing. We performed the same preprocessing procedure on both the hue- and motion- scaling datasets. We first excluded outlying trials by stimulus angle for each observer. We used a jackknife technique to detect outliers. This technique consisted of computing the mean and SD for each possible group of three trials. We then determined whether for the closest cluster of three trials (i.e., the cluster with the smallest SD) the fourth trial was more than 3 SDs away from the mean of that cluster. In that case, the fourth trial would be excluded from subsequent analyses. Otherwise, all trials were included. Of the total trials, 14% were excluded by this criterion for both the hue scaling and motion scaling. We then applied criteria for excluding observers with responses across trials deemed inconsistent. We determined the threshold for exclusion by calculating each observer's average response SD across trials and excluding the full dataset for any observer whose average SD across trials exceeded the average SD across trials for the group by 3 SDs. By this criterion, none of the observers' hue-scaling functions were excluded, while one observer's motion-scaling function was excluded. Lastly, given that subsequent analyses were performed on the mean scaling functions (i.e., averaging across trials for each observer), we excluded outlying points across the group by each stimulus angle. We determined the threshold of exclusion at each stimulus angle as any response that was more than 4 SDs away from the mean response at each stimulus angle. We used a more inclusive criterion for this case to avoid missing data. By this criterion, 0.5% of points were excluded from the group of mean hue-scaling functions, and 0.7% of points were excluded from the group of mean motion-scaling functions. Imputed values were determined for the missing values to eliminate the need for smoothing of the data correlation matrix for factor analysis, which would be required if the data were instead missing due to pairwise deletion. We used a nonparametric missing value imputation by random forest (48) for the hue-scaling data, given its deviation from multivariate normality. This imputation was performed with the "missForest" package in R using the missForest function. Since the motion-scaling data met assumptions of multivariate normality, we instead used multivariate imputation by chained equations (49). This imputation was performed with the "mice" package in R using the mice function with the predictive mean-matching method. All other analyses were performed on the hue- and motion- scaling datasets following these preprocessing steps.

Arcsine Transform. For analyses where indicated, this transform was applied to equate the variances for different mean percentages. The specific form used was:

$$R_{arc} = \frac{(2(\sin^{-1}(\sqrt{R}))}{\pi},$$

R corresponds with the percentage the observer provided in the hue-scaling task, and R_{arc} the arcsine-transformed percentage. The limits of the scale values transformed in this way are still 0 and 100. This transform is commonly used on hue-scaling data (46) and is preferred to logit given that values close to 0% and 100% are common in hue-scaling responses.

Test–Retest and Split-Half Reliabilities. Observers' scaling functions in each session were first detrended by subtracting the mean scaling function across all observers for that session. To assess test–retest reliability, the detrended responses across sessions were correlated by taking the average scaling function for each observer for sessions 1 and 2, concatenating all observer scaling functions for session 1 and also for session 2, and calculating a correlation between sessions. To assess split-half reliability, detrended responses were concatenated for the first trials of each session and for the second trials of each session. A correlation between first and second trials within each session was then calculated to determine split-half reliability.

Factor Analysis. The same statistical analyses were performed on both hue-scaling and motion-scaling data sets. Each set consisted of a mean hue-scaling function (averaging across trials) for the 46 and 48 included observers for the hue and motion sets, respectively. Factor analyses were computed on the response correlation matrices in each case (Fig. 2*C*). For the principal component (PC) method of estimation, the factor analysis was computed using the principal function within the R CRAN Package "psych." The appropriate number of factors and rotation method were specified for the function to calculate the corresponding factor model. For each dataset, we confirmed the existence of a correlation structure, and therefore that factor analysis was appropriate (Bartlett's test of sphericity; color: $\chi^2(630) = 1{,}246.00$, $P < 0.001$; motion: $\chi^2(630) = 1336.97$, $P < 0.001$). We then decided on the method of estimation, rotation, and number of factors for each solution using the following criteria. To choose the appropriate method of estimation, we tested whether each dataset met assumptions of multivariate normality. Because the hue-scaling dataset deviated from multivariate normality (i.e., the statistical distances were not chi-square distributed according to a one-sample Kolmogorov–Smirnov test ($D = 0.21$, $P = 0.03$)), we extracted the factors using a principal component method as it is robust to normality deviations unlike other common estimation methods (e.g., maximum likelihood) (28). However, a maximum likelihood method yielded similar results (*SI Appendix*, Fig. S1). Though the motion-scaling dataset met assumptions of multivariate normality (i.e., the statistical distances were chi-square distributed according to a one-sample Kolmogorov–Smirnov test ($D = 0.17$, $P = 0.11$)), we also extracted the factors using the principal component method for consistency with the factor model reported for color. A maximum likelihood method again yielded similar results (*SI Appendix*, Fig. S4). After extraction, we applied a Varimax rotation to the estimated factors for each dataset (an orthogonal rotation that favors a sparse factor structure) because an oblique solution (allowing for the factors to be correlated), yielded similar results. Specifically, the Spearman's rank-order correlation between the oblique (Promax) and orthogonal (Varimax) factor patterns was significant for both the hue-scaling ($\rho(358) = 0.49$, $P < 0.001$) and motion-scaling ($\rho(322) = 0.94$, $P < 0.001$) datasets.

Systematic Loadings. Finally, we chose the number of factors using a "systematic loadings" criterion, which distinguishes significant factors from "noise", as those with high loadings (i.e., the amount of variance in the variable that is accounted for by the factor) on multiple, consecutive stimulus angles. This criterion has been useful for identifying meaningful factors within stimulus spaces that lie on a continuum, particularly for color (50). For the hue-scaling data, factors with at least two consecutive loadings higher than 0.63 were determined as systematic using a bootstrapping procedure, as they were unlikely to have emerged by chance ($P = 0.014$). For the motion-scaling data, factors with at least two consecutive loadings higher than 0.68 were determined as systematic for these data using a bootstrapping procedure, as they were unlikely to have emerged by chance ($P = 0.0003$).

Systematic factors were determined using a custom-built algorithm that searches for a solution where all factors included are systematic (i.e., with multiple high, consecutive loadings). Bootstrapping procedures were used to determine the threshold that constituted a "high" loading (loading threshold) and the threshold number of "consecutive" loadings (run threshold) that were required for a factor to be considered systematic, given the parameters of the dataset (e.g., number of observers, number of variables, type of model). The bootstrapping analysis to determine the loading and run thresholds is accomplished by factor-analyzing correlation matrices of random datasets (of the same dimensions as the original) with the same method of estimation (e.g., principal components) and rotation (e.g., Varimax) that was applied to the original data. From this, it is possible to determine whether a factor with a given loading and run threshold could have emerged by chance rather than corresponding to a meaningful, latent source of variability (i.e., systematic). The final factor analysis solutions required that all included factors were systematic according to these criteria, and the solution with the highest number of factors where all were systematic was chosen for each dataset. For each dataset, the algorithm starts by extracting as many factors as there are variables, and iteratively decreases the number of extracted factors until all factors in the solution are systematic. For each number of factors, the algorithm begins by searching for a solution with the consecutive loading threshold set at two, and iteratively increases the consecutive loading threshold up to a limit, again while searching for a solution where all factors are systematic. For each solution, the algorithm finds the highest consecutive loadings for each factor, and the bootstrapping analysis determines the frequency that a solution emerges by chance with the given loading threshold (set as the minimum loading of the maximum consecutive loadings across all extracted factors) and number of consecutive loadings. It then searches for a solution where the frequency is lower than 5% (i.e., $P < 0.05$) by adding to the consecutive loading threshold up to a limit and decreasing extracted factors until this criterion is met. Therefore, the loading and run thresholds to classify factors as systematic were slightly different for each solution. The bootstrapping analysis to determine the threshold values was performed with custom code written in R (https://github.com/kjemery/factor_analysis).

Factor Scores. We estimated the factor scores by Thurstone's (29) least-squares regression approach, which maximally ensures that the factor scores are correlated to the estimated factor (29, 51). Factor scores were calculated with the "psych" package by using the factor.scores function and the "Thurstone" method.

Code Availability. All analyses were conducted using both available and custom code written in MATLAB and R, which can be found here: https://github.com/kjemery/factor_analysis.

Data, Materials, and Software Availability. The data (52) that support the findings in this study can be found here: https://github.com/kjemery/factor_analysis.

ACKNOWLEDGMENTS. Research funded by National Eye Institute (EY010834).

1. S. K. Shevell, F. A. Kingdom, Color in complex scenes. *Annu. Rev. Psychol.* **59**, 143–166 (2008).
2. E. Hering, *Zur Lehre vom Lichtsinne: Sechs Mittheilungen an die Kaiserl. Akademie der Wissenschaften in Wien* (C. Gerold's Sohn, 1878).
3. L. M. Hurvich, D. Jameson, An opponent-process theory of color vision. *Psychol. Rev.* **64**, 384 (1957).
4. A. M. Derrington, J. Krauskopf, P. Lennie, Chromatic mechanisms in lateral geniculate nucleus of macaque. *J. Physiol.* **357**, 241 (1984).
5. D. M. Dacey, Parallel pathways for spectral coding in primate retina. *Annu. Rev. Neurosci.* **23**, 743–775 (2000).
6. R. Shapley, M. J. Hawken, Color in the cortex: Single-and double-opponent cells. *Vision Res.* **51**, 701–717 (2011).
7. Y. Xiao, Y. Wang, D. J. Felleman, A spatially organized representation of colour in macaque cortical area V2. *Nature* **421**, 535–539 (2003).
8. B. R. Conway, S. Moeller, D. Y. Tsao, Specialized color modules in macaque extrastriate cortex. *Neuron* **56**, 560–573 (2007).
9. M. Li, F. Liu, M. Juusola, S. Tang, Perceptual color map in macaque visual area V4. *J. Neurosci.* **34**, 202–217 (2014).
10. K. S. Bohon, K. L. Hermann, T. Hansen, B. R. Conway, Representation of perceptual color space in macaque posterior inferior temporal cortex (the V4 complex). *ENeuro* **3**, 28 (2016).
11. K. Koida, H. Komatsu, Effects of task demands on the responses of color-selective neurons in the inferior temporal cortex. *Nat. Neurosci.* **10**, 108–116 (2007).
12. G. J. Brouwer, D. J. Heeger, Decoding and reconstructing color from responses in human visual cortex. *J. Neurosci.* **29**, 13992–14003 (2009).
13. I. Kuriki, P. Sun, K. Ueno, K. Tanaka, K. Cheng, Hue selectivity in human visual cortex revealed by functional magnetic resonance imaging. *Cerebral Cortex* **25**, 4869–4884 (2015).
14. R. L. De Valois, I. Abramov, G. H. Jacobs, Analysis of response patterns of LGN cells. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **56**, 966–977 (1966).
15. M. A. Webster, The verriest lecture: Adventures in blue and yellow. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **37**, V1–V14 (2020).
16. G. Malkoc, P. Kay, M. A. Webster, Variations in normal color vision. IV. Binary hues and hue scaling. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **22**, 2154–2168 (2005).
17. K. J. Emery, V. J. Volbrecht, D. H. Peterzell, M. A. Webster, Variations in normal color vision. VI. Factors underlying individual differences in hue scaling and their implications for models of color appearance. *Vision Res.* **141**, 51–65 (2017).

18. J. Bosten, A. Lawrance-Owen, No difference in variability of unique hue selections and binary hue selections. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **31**, A357–A364 (2014).
19. J. Bosten, A. Boehm, Empirical evidence for unique hues? *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **31**, A385–A393 (2014).
20. S. K. Shevell, P. R. Martin, Color opponency: Tutorial. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **34**, 1099–1108 (2017).
21. C. Matera *et al.*, A comparison of two methods of hue scaling. *J. Vision* **19**, 298b–298b (2019).
22. J. D. Mollon, J. M. Bosten, D. H. Peterzell, M. A. Webster, Individual differences in visual science: What can be learned and what is good experimental practice? *Vision Res.* **141**, 4–15 (2017).
23. D. H. Peterzell, J. F. Kennedy, Discovering sensory processes using individual differences: A review and factor analytic manifesto. *Electron. Imaging* **2016**, 1–11 (2016).
24. M. Livingstone, D. Hubel, Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science* **240**, 740–749 (1988).
25. L. Ungerleider, J. V. Haxby, 'What'and 'where' in the human brain. *Curr. Opin. Neurobiol.* **4**, 157–165 (1994).
26. M. A. Goodale, A. D. Milner, Separate visual pathways for perception and action. *Trends Neurosci.* **15**, 20–25 (1992).
27. D. Jameson, L. M. Hurvich, Perceived color and its dependence on focal, surrounding, and preceding stimulus variables. *JOSA* **49**, 890–898 (1959).
28. R. A. Johnson, D. W. Wichern, "Factor analysis and inference for structured covariance matrices" in *Applied Multivariate Statistical Analysis* (Pearson, 2014), pp. 378–430.
29. L. L. Thurstone, *The Vectors of Mind: Multiple-Factor Analysis for the Isolation of Primary Traits* (University of Chicago Press, 1935).
30. T. Regier, P. Kay, R. S. Cook, Focal colors are universal after all. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 8386–8391 (2005).
31. C. Witzel, K. R. Gegenfurtner, Color perception: Objects, constancy, and categories. *Annu. Rev. Vision Sci.* **4**, 475–499 (2018).
32. M. A. Webster, P. Kay, Color categories and color appearance. *Cognition* **122**, 375–392 (2012).
33. B. C. Hansen, E. A. Essock, A horizontal bias in human visual processing of orientation and its correspondence to the structural components of natural scenes. *J. Vision* **4**, 5–15 (2004).
34. T. Namima, M. Yasuda, T. Banno, G. Okazawa, H. Komatsu, Effects of luminance contrast on the color selectivity of neurons in the macaque area v4 and inferior temporal cortex. *J. Neurosci.* **34**, 14934–14947 (2014).
35. S. L. Buck, T. DeLawyer, Dark versus bright equilibrium hues: Rod and cone biases. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **31**, A75–A81 (2014).
36. A. Anzai, G. C. DeAngelis, Neural computations underlying depth perception. *Curr. Opin. Neurobiol.* **20**, 367–375 (2010).
37. M. Wexler, J. J. Van Boxtel, Depth perception by the active observer. *Trends Cognit. Sci.* **9**, 431–438 (2005).
38. P. Lennie, J. Krauskopf, G. Sclar, Chromatic mechanisms in striate cortex of macaque. *J. Neurosci.* **10**, 649–669 (1990).
39. R. L. De Valois, E. W. Yund, N. Hepler, The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res.* **22**, 531–544 (1982).
40. Q. Zaidi, B. Conway, Steps towards neural decoding of colors. *Curr. Opin. Behav. Sci.* **30**, 169–177 (2019).
41. B. Berlin, P. Kay, *Basic Color Terms: Their Universality and Evolution* (Univ of California Press, 1991).
42. T. J. Andrews, D. Schluppeck, Ambiguity in the perception of moving stimuli is resolved in favour of the cardinal axes. *Vision Res.* **40**, 3485–3493 (2000).
43. M. C. Morrone, D. C. Burr, S. Di Pietro, M. A. Stefanelli, Cardinal directions for visual optic flow. *Curr. Biol.* **9**, 763–766 (1999).
44. D. I. A. MacLeod, R. M. Boynton, Chromaticity diagram showing cone excitation by stimuli of equal luminance. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **69**, 1183–1186 (1979).
45. M. A. Webster, E. Miyahara, G. Malkoc, V. E. Raker, Variations in normal color vision. II. Unique hues. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **17**, 1545–1555 (2000).
46. I. Abramov, J. Gordon, H. Chan, Color appearance in the peripheral retina: Effects of stimulus size. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **8**, 404–414 (1991).
47. J. Gordon, I. Abramov, H. Chan, Describing color appearance: Hue and saturation scaling. *Percept. Psychophys.* **56**, 27–41 (1994).
48. D. J. Stekhoven, P. Bühlmann, MissForest–non-parametric missing value imputation for mixed-type data. *Bioinformatics* **28**, 112–118 (2012).
49. S. van Buuren, K. Groothuis-Oudshoorn, Mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software* **45**, 1–67 (2015).
50. M. A. Webster, D. I. A. MacLeod, Factors underlying individual differences in the color matches of normal observers. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **5**, 1722–1735 (1988).
51. C. DiStefano, M. Zhu, D. Mindrila, Understanding and using factor scores: Considerations for the applied researcher. *PARE* **14**, 20 (2009).
52. K. J. Emery, V. J. Volbrecht, D. H. Peterzell, M. A. Webster, Factor analysis protocol. *factor-analysis*. https://github.com/kjemery/factor_analysis, April 17, 2020.