Temporal dynamics of the neural representation of hue and luminance polarity

Katherine L. Hermann[1,4,*], Shridhar R. Singh[1*], Isabelle A. Rosenthal[1,5,*], Dimitrios Pantazis[2], Bevil R. Conway[1,3,†]

1. Laboratory of sensorimotor research, National Eye Institute, Bethesda MD 20892
2. McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge MA 02139
3. National Institute of Mental Health, Bethesda MD 20892
4. Current address: Department of Psychology, Stanford University, Stanford CA 94305
5. Current address: Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, California 91125

*equal contribution

[†]Corresponding author: bevil@nih.gov

Author contributions: KH and BRC conceived and designed the experiments; KH and IR collected the data; KH, IR, and SS analyzed the data; IR, SS and BRC made the figures; DP provided expertise and resources for MEG; BRC supervised the work and wrote the paper.

Competing Interests: The authors declare no competing interests.

Acknowledgements: We thank Chris Baker, Susan Wardle, Santani Teng, and Qasim Zaidi for helpful discussions, and Stuart Duffield for help checking and documenting the analysis pipeline. This research was supported in part by the Intramural Research Program of the National Institutes of Health, National Eye Institute to BRC.

Abstract

Hue and luminance contrast are basic visual features, yet the timing of the neural computations that extract them, and whether they depend on common neural circuits, is not well established. Using multivariate analyses of magnetoencephalography data, we show that hue and luminance-contrast polarity can be decoded from MEG data and, with lower accuracy, both features can be decoded across changes in the other feature. These results are consistent with the existence of both common and separable neural mechanisms. The decoding time course is earlier and more temporally precise for luminance polarity than hue, a result that does not appear to depend on task, suggesting that luminance contrast is an updating signal that separates visual events. Meanwhile, cross-temporal generalization is slightly greater for representations of hue compared to luminance polarity, providing a neural correlate of the preeminence of hue in perceptual grouping and memory. Finally, decoding of luminance polarity varies depending on the hues used to obtain training and testing data; the pattern of results suggests that luminance contrast is mediated by both L-M and S cone sub-cortical mechanisms.

Introduction

The most basic computations performed by the visual system yield hue and luminance contrast. *Hue* is that property of color referred to by a color name (e.g. “blue”, “green”). *Luminance contrast* is related to how light or dark a color is, which is, importantly, distinct from absolute luminance. Although hue and luminance contrast are often depicted as independent dimensions,[1] the relationship between these features is not well understood[2, 3], and the extent to which luminance is carried by both cone-opponent retinal mechanisms (L-M and S) remains unclear[4, 5, 6].

On the one hand, the demonstration of independent frequency shifts for color and luminance[7] and the selective impact of adaptation on detection thresholds[8, 9] suggests that hue and luminance contrast are processed by separate neural channels, which may be encoded by the separate magnocellular and parvocellular channels in the lateral geniculate nucleus[10]. Separate encoding is also evident in convolutional neural networks trained for object recognition, which show independent filters for hue and luminance contrast in the earliest layer[11, 12, 13, 14]. Color and luminance edges are also relatively independent in natural scene statistics[15]. Moreover, hue and luminance contrast are distinguished by their efficiency in visual memory tasks: short-term visual memory is better for color than for luminance contrast when stimuli are matched in cone contrast.[16]

On the other hand, the many perceptual interactions of hue and luminance contrast suggest these features are processed together. These interactions are evident in the perception of orientation,[17] judgments of brightness across colors[18], color categorization[19], color naming[20, 21, 22], and color decoding with

magnetoencephalography[23]. In addition, hue can mask changes in luminance contrast[24]. Consider 8 colored spirals: four hues of both light and dark luminance polarity (Figure 1a). Most people will group them by hue (rows), not the sign of luminance contrast (columns). Taken together, these data underscore long-standing questions about the extent to which hue and luminance contrast are encoded by the same or separate neural mechanisms and the timing of these operations.

A prominent theory, which reflects the interactions of hue and luminance contrast evident in perception, is that within visual cortex, hue is processed not in isolation, but together with information about luminance contrast[26]. This theory is supported by the observation that most V1 cells receive a mixture of inputs from subcortical channels and respond to both color and luminance contrast[27, 28, 29]. But a small population of V1 cells show striking responses to color stimuli that lack luminance contrast[29, 30, 31, 32].Thus the V1 neurophysiology cannot rule out the possibility that hue and luminance contrast could be encoded by parallel pathways[33]. which may contribute differentially to processing in parallel routes through extrastriate areas[34, 35, 36, 37, 38]. If hue and luminance contrast are encoded separately, one might expect luminance contrast to be computed earlier than hue because magnocellular neurons have shorter latencies than parvocellular neurons. But because there are relatively fewer magnocellular neurons, their latency advantage may be lost through convergence in visual cortex[25]. The time taken by the brain to encode hue and luminance contrast remains poorly understood.

Clues to how color and luminance contrast are processed by the brain have been provided by univariate visual evoked potential measurements to equiluminant and achromatic stimuli[39, 40]. But it has not been possible to infer from these experiments the underlying neural mechanisms because all subcortical channels respond to nominally equiluminant stimuli[41]. Moreover, such experiments are inconclusive about timing because response latency depends on stimulus contrast, and there is no single method for equating color contrast and luminance contrast[42]. Clues have also been provided by fMRI[43, 44, 45], but fMRI does not predict results from behavioral adaptation[9], and cannot uncover the relevant temporal dynamics because it is limited by the relatively sluggish time course of blood flow[46].

In this work, we explore another approach to address the neural mechanisms for color, analyzing the MEGco dataset[23] (MEG color decoding dataset) using multivariate analysis to isolate common and separable representations for hue and luminance contrast. The MEGco dataset was obtained to address

two sets of questions. The first set of questions, addressed by Rosenthal et al.[23], concerns the similarity relationships of the patterns of neural activity elicited by different colors. That analysis shows that hue and the sign of luminance contrast (light versus dark) interact in the representation of color, and these interactions are consistent with color-naming patterns: hues labeled with different names for light and dark versions (e.g. yellow/brown) elicit less similar patterns of MEG activity for the light and dark versions, compared to hues that are given the same name across lightness (e.g. green). The second set of questions, addressed presently, concerns the extent to which representations of hue and luminance polarity are separable, and importantly, the temporal dynamics of the representations for hue and luminance polarity. Can hue be decoded independently of luminance polarity, and if so with what timing? Can luminance polarity be decoded independently of hue, and if so with what timing? To what extent do separable representations of hue and luminance polarity generalize across time? Although these questions are related to ones addressed in our previous report, they are distinct and cannot be answered using previously published analyses. The results here show that hue and luminance polarity can be decoded independently from MEG activity, with hue reaching peak decoding about 20 ms after luminance polarity and having slightly greater cross-temporal decodability. We also report here source localization of the MEG data to functionally defined regions identified with functional magnetic resonance imaging (fMRI) in the same participants in whom the MEG data were obtained.

Results

The experiments were designed to enable a decoding analysis to answer two over-arching questions. Given patterns of MEG activity, is it possible to independently decode the hue and luminance polarity (light or dark) of the stimulus that elicited the activity? And what is the time course – that is, how long does it take the brain to extract these features? We addressed these questions by analyzing MEG responses in 18 participants, obtained while the participants were shown spirals that could appear in one of eight colors that varied in hue and were either a luminance increment or decrement (Figure 1).

The stimuli were suprathreshold and roughly matched in absolute color contrast and luminance contrast, in terms of detection-threshold units (contrasts were >30x detection threshold, see Methods). Equating stimuli in units of detection threshold is one method among many for comparing color and luminance responses[47, 48, 49, 50]. The challenge in comparing responses to color and luminance stimuli arises in part because response magnitude can depend on stimulus contrast[39], especially for near-threshold stimuli.

Multivariate analyses of suprathreshold stimuli ameliorate this problem because they uncover any difference in the pattern of response between a pair of variables, providing insight about the information processed by the brain[51].

The colors were defined by the DKL color space[52, 53], which is constructed from the cone-opponent dimensions that correspond to retinal color-encoding mechanisms (Supplementary Figure 1). The stimuli were set to the intermediate directions in the space; thus all stimuli involved modulations of both L-M cone activity and S-cone activity, and all stimuli had the same magnitude of modulation of L-M, and the same magnitude of modulation of S. Different colors were created by pairing different signs of these modulations. The luminance contrast of all the stimuli was fixed (25%) but varied in polarity (light or dark) relative to the adapting background. If hue and the polarity of luminance contrast are encoded by separable neural mechanisms, it should be possible to decode hue even if the MEG data used to train the classifiers were elicited by stimuli that differed in luminance polarity from the test stimuli; and it should be possible to decode luminance polarity even if the classifiers were trained using data elicited by stimuli that differed in hue from the test data. In other words, the results should show evidence that hue decoding generalizes across luminance polarity; and luminance-polarity decoding generalizes across hue. The time course should tell us about the relative stage in the visual-processing hierarchy at which hue and luminance polarity representations are encoded and/or the relative amount of recurrent processing required for each computation. Alternatively, if hue and luminance polarity are encoded together, it should be possible to decode specific hue-luminance combinations, but not each dimension invariant to the other dimension. Given these alternatives, it was important for the experiment to have enough power, which we ensured by first conducting an extensive pilot experiment to determine the number of trials needed to obtain reliable data (Supplementary Figure 2).

We used a maximum correlation coefficient classifier (as implemented in the Neural Decoding Toolbox[54], see Methods). We performed within-participant decoding (trained and evaluated classifiers on data in each participant), independently for each timepoint (applied independent classifiers at each time point relative to stimulus onset). All analyses involved pair-wise comparisons (Figure 1) and were cross validated (Figure 2), yielding plots that show how representations unfold over time. Participants were told to maintain fixation throughout stimulus presentation and to blink at designated times. Data during eye blinks or breaks in fixation were removed (see Methods). To control fixation and attentional state, participants

engaged in a 1-back hue-matching task: every 3-5 trials, the participants were queried with a "?" on the screen to report via button press whether the two preceding stimuli matched in hue. The impact of task was assessed in a control experiment, in which two participants performed half the trials using a 1-back hue matching task and the other half of the trials using a 1-back luminance-contrast matching task. Task had no impact on the main conclusions (see Figure 5; see Supplementary Figure 5 for behavioral results). We used a spiral-shaped stimulus to avoid cardinal or radial response biases[55, 56, 57].

The eight colors ("conditions"; Figure 1) were presented in pseudo-random order for 116 ms with 1 second of the gray background between presentations (1 sec inter-stimulus interval). We collected responses to a very large number of trials of each condition (N=500), removed trials with artifacts such as eye blinks, and randomly subsampled the remaining trials to obtain 375 trials per condition. Figure 2 illustrates an analysis in which classifiers were trained to decode luminance polarity given patterns of MEG data elicited by bright and dark pink (the decoding curve bottom right is real data, the rest of the figure is for illustration purposes). The classifiers were tested on separate data elicited by the same stimuli, bright and dark pink. The results reveal the classification accuracy for luminance polarity carried by a specific hue (pink).

The decoding problem shown in Figure 2 is referred to as a luminance-polarity identity problem because the hue of the stimuli from trials used to train the classifier is identical to the hue of the stimuli in the test trials. Figure 1a, left, shows schematically the full set of luminance-polarity identity decoding problems: the stimuli associated with the training of classifiers are indicated by solid arrows, and the stimuli associated with the corresponding tests are shown by dashed arrows; shading of the arrows links the training and testing conditions for a given problem. In other problems of luminance-polarity decoding, the hue of the stimuli differed between trials used to train versus test the classifier. For example, classifiers were trained using patterns of MEG activity elicited by light and dark pink but tested using activity elicited by light and dark blue, or light and dark orange, or light and dark green. We refer to these problems as generalizing-across-hues since they uncover the extent to which luminance polarity can be decoded independent of hue (Figure 1a, right).

In other analyses we determined the extent to which classifiers could decode hue identity (Figure 1b, left), and hue generalizing-across-luminance-polarity (Figure 1b, right). The generalization problems provide a test of invariance: decoding luminance-polarity invariant to hue, and decoding hue invariant to luminance

polarity. Supplementary Figure 3 shows the test-retest reliability of the data from the main experiments[58]. All analyses involved binary classifiers (chance is 50%) to facilitate a direct comparison of the results among the different problems.

Decoding luminance polarity

Figure 3a shows the average classifier performance across the four luminance-polarity identity problems (solid line) and the 12 generalizing-across-hues problems (dashed line; shading around each trace shows the bootstrapped standard error). The results show that classification accuracy was significantly above chance for luminance polarity for both types of problems. Thus it was possible to decode luminance polarity independent of hue, which supports the hypothesis that the brain has a representation of luminance contrast that is independent of the representation of hue. The magnitude of decoding does not provide a measure of the absolute size of an effect, but it is nonetheless a valid measure of relative effect sizes within a given study[51]. Peak decoding accuracy was higher for the identity problems compared to the generalization problems (84% [81, 87] versus (69% [65, 74]; square brackets contain the 95% C.I obtained by bootstrapping). This result supports the hypothesis that the brain not only has a representation of luminance polarity that is generalized across hues, but also has a representation of luminance polarity that is combined with the representation of hue. In other words, these results show that the brain has a representation of luminance polarity carried by channels that are hue selective.

Let us compare the time course of decoding for the identity problems and the generalization problems. The latency at which decoding became significant was not different for the identity problems (65ms [60, 80]) versus the generalization problems (55 ms [50, 70], based on bootstrapping, n=18, p=0.0.058). The time to peak was not different for the identity problems compared to the generalization problems (inset Figure 3a, p=0.34; dashed vertical line, 96.5 ms [85, 100], versus solid vertical line, 104.4 ms [100, 110]; confidence limits computed by 1000 bootstrap draws across participants). Close inspection of the data shows a trend in favor of an earlier onset time and earlier time to peak for the generalization problems (generalization problems showed equal or earlier time to peak than identity problems, p=0.002); this observation implies that the time-to-peak is a reliable measure of how much time the brain takes to generate a representation because it does not vary in a trivial way with changes in the amplitude of decoding (one might have expected time-to-peak to be longer for curves with lower amplitude decoding; the data show, if anything,

222 the opposite: the problem with lower amplitude decoding achieves significance slightly earlier). Following
223 peak decoding, the generalization problems showed a pronounced dip. Finally, both the identity and
224 generalization problems had a second prominent decoding peak (open arrowhead) following the initial
225 peak (curved arrow; to facilitate comparison, the y-axis value of the open arrowhead is the height of the
226 initial peak of the identity decoding problems). For the generalization problem, the second peak had the
227 same or greater amplitude as the initial peak. We attribute the first peak to decoding the onset of the
228 stimulus and the second peak to decoding the cessation of the stimulus.
229
230 *Decoding hue*
231
232 Figure 3b shows the average performance across the 12 hue identity problems (solid line) and the 12
233 generalizing-across-luminance-polarity problems (dashed line). The plot shows that hue was decodable in
234 both cases. That it was possible to decode hue generalizing across luminance polarity supports the
235 hypothesis that the brain has a representation of hue that is separate from the representation of luminance
236 contrast. The latency at which decoding became significant was 115 ms [35, 355] for the generalization
237 problems, and 80 ms [75, 85] for the identity problems (the latency of the generalization problems was
238 greater than the upper 95% CI limit for the latency of the identity problems, but a direct comparison of the
239 bootstrapped values was not significant, p=0.09). As with the luminance-polarity decoding problems,
240 decoding had a higher peak magnitude for the identity problems (78% [74, 82]) compared to the
241 generalization problems (56% [53, 59]), which provides support for the hypothesis that the brain also has a
242 representation of hue that is inseparable from the representation of luminance contrast. The time to peak
243 was not different for the identity and generalization problems (solid vertical line, 121.8 ms [120, 130];
244 dashed vertical line, 119.1 ms [115, 130]; p=0.82). Moreover, the time to peak was not significantly
245 different for identity problems among luminance-contrast increments (122 ms [115, 130]) compared to
246 identity problems among luminance-contrast decrements (123 ms [115, 130]); and the magnitude of peak
247 decoding accuracy for these two sets of problems also were not significantly different (for light stimuli: 77%
248 [73, 82]; for dark stimuli: 78% [74, 83]) (Supplementary Figure 4). These results show that any potential
249 difference in saturation between the light and dark stimuli does not influence the time course of decoding
250 (saturation is ill-defined, but as discussed in the Methods, there is an argument that the saturation of the
251 dark stimuli was higher than the saturation of the light stimuli). The results of Supplementary Figure 4 also
252 support the argument that the differences in timing of decoding hue and luminance polarity cannot be

attributed to subtle differences in the contrast of the hue and luminance-polarity stimuli. Finally, to the extent the solid line in Figure 3b shows a second peak (open arrowhead), it was less pronounced than the initial decoding peak (curved arrow) (as in Figure 3a, the y-axis value of the open arrowhead is at the height of the initial peak for the identity decoding problem).

Comparing the temporal dynamics of decoding hue and luminance polarity

To directly compare the timing of hue and luminance-polarity problems, we subsampled the data to obtain equal peak decoding across the two sets of problems[59] (the height of the initial decoding peaks of the blue and black traces are matched in Figure 4a and in Figure 4b). The times to peak for the subsampled luminance-polarity data were 99.8 ms [95, 110] for the generalization problem and 103.4 ms [100, 110] for the identity problem. The comparison of the hue and luminance-polarity identity problems (Figure 4a) and of the hue and luminance-polarity generalization problems (Figure 4b) underscores three differences in time course of decoding. First, hue was decodable after luminance polarity as assessed by the time of peak decoding (identity problems: 18 ms delay, $p=0.003$; generalization problems: 19 ms delay, $p=0.017$). Differences in the latency of decoding onset showed a trend that reflects these timing differences, but the differences in latency of decoding onset were not significant (for identity problems, $p=0.13$; for generalization problems, $p=0.09$). Second, the time of peak hue decoding corresponded to the dip in the luminance-polarity decoding curve, and vice versa, the time point of peak luminance-polarity decoding corresponded to a notch in the hue-decoding curve. This was especially clear in the generalization curves, where the black and blue traces are in counterphase (quantified as the Spearman correlation of the derivative of the decoding curves, computed for 116 ms following the average of the onset latencies of the hue and luminance-polarity generalization problems; $R=-0.53$, $p=0.01$). Third, for the identity problems, the second peak (open arrowhead) was larger for luminance polarity than for hue (double-headed arrow, $p<0.001$; note that precise p values are stated throughout the paper except when the p value must be smaller than $p=0.001$. In those cases, because the p values were derived from 1000 bootstrap iterations, the precise p value is unknown and the p value is listed as <0.001). The same trend was observed for the generalization problems, but it was not significant ($p=0.3$).

Impact of task on decoding

Decoding accuracy of visual stimuli from patterns of MEG activity could be impacted by task engagement during MEG data collection[60, 61]. Task effects can enhance task-relevant object features and are reported to arise relatively late after stimulus onset[62, 63]; nonetheless, could the differences in timing for decoding hue versus luminance polarity be attributed to the task that participants performed during MEG data collection? We answered this question in a control experiment conducted prior to the main experiments. The control experiment was identical to the main experiments except that for half the runs the participants performed a 1-back hue-matching task, and for the other half of the runs the participants performed a 1-back luminance-polarity matching task (Figure 5a; each participant completed 5 sessions, one of which was not analyzed because of a technical error in data acquisition; each trial involved a 116 ms stimulus presentation and a 1 sec ISI; the task alternated between runs in each session). Participants learned both tasks, achieving excellent accuracy within the first MEG data-acquisition session and improving in reaction time over consecutive sessions (Supplementary Figure 5).

The time to peak decoding was quantified by training and testing classifiers for each session and each participant (4 sessions x 2 participants for each task type) and bootstrapping across the eight curves for a given task. Peak decoding times were different for hue than for luminance polarity, for data collected using both tasks: decoding of hue peaked later than decoding of luminance polarity (Figure 5b; difference in time to peak for data obtained using the 1-back luminance task, $p<0.001$; difference in time to peak for data obtained using the 1-back hue task, $p<0.001$). The time to peak for decoding hue was 112 ms [106, 119] using data obtained with the hue task, and 111 ms [101, 121] using data obtained with the luminance task; while the time to peak for decoding luminance polarity was 89 ms [84, 95] for the hue task, and 89 ms [81, 101] for the luminance task. These results replicate the results from the main experiments, showing that decoding luminance polarity reached a peak about 20 ms earlier than decoding hue. The time to peak decoding of luminance polarity was indistinguishable for data collected under the two task conditions (Figure 5c, left panel; $p=0.384$). Similarly, the time to peak decoding of hue was indistinguishable for data collected under the two task conditions (Figure 5c, right panel; $p=0.427$). These results show that task did not impact the time to peak, for decoding hue or luminance polarity. The data collected in the control experiment also show that the relative timing difference for decoding hue versus luminance polarity was apparent for each individual subject: combining data from both tasks, the time to peak for decoding hue for subject 1 was 108 ms [101, 117], and for subject 2 was 114 ms [110, 118]); the time to peak for decoding luminance polarity for subject 1 was 88 ms [86, 90], and for subject 2 was 92 ms [84, 102]). The time to

peak for decoding hue was different from the time to peak for decoding luminance polarity in both subjects (both S1 and S2, p=0.001).

The results of the control experiment provide evidence of one possible impact of task: the duration of significant decoding. Decoding of luminance polarity was significant for a longer duration when participants were engaged in the 1-back luminance task compared to when they were engaged in the 1-back hue task (Figure 5c, left panel, horizontal lines below the decoding trace show time points of significant decoding). Similarly, decoding of hue was significant for a longer duration when participants were engaged in the 1-back hue task compared to when they were engaged in the 1-back luminance task (Figure 5c, right panel). The results of the control experiment are consistent with the manifestation of task-related effects arising relatively late (>300 ms after stimulus onset), as documented by others[62], and argue against the hypothesis that the difference in time-to-peak for decoding hue and luminance polarity, which is evident much earlier than 300 ms, be attributed to task.

Greater cross-temporal decoding for hue than luminance polarity

The results discussed so far evaluate the classifiers' performance using test data obtained at the same time point after stimulus onset as the data used to train the classifier. The classifiers show significant decoding for a substantial amount of time, hundreds of milliseconds, following stimulus onset. One possibility is that the pattern of activity is relativity sustained over this time period; another possibility is that it is dynamic and transient[64]. To distinguish between these alternatives, we trained classifiers on the patterns of activity at each point in time and evaluated the extent to which they could predict activity at all other time points. If activity patterns are dynamic, the cross-temporal analysis will recover strong decoding performance only for situations in which the training and testing data sets were obtained at the same timepoint relative to stimulus onset, i.e. along the diagonal in a cross-temporal decoding plot. Alternatively, if activity patterns are more sustained, the analysis will show strong decoding performance at time points away from the diagonal.

Figure 6 shows the cross-temporal plots for decoding luminance polarity generalizing across hues (Figure 6a, left), and for decoding hue generalizing across luminance polarity (Figure 6b, right); the figures show the results averaged over the 12 individual problems for luminance polarity and hue using the subsampled

data shown in Figure 4, which ensures that any differences in cross-temporal decoding cannot be not attributed to absolute differences in peak decoding. The color scale in the heatmap shows the percent classification—the values along the diagonal are the same as those shown in Figure 4b. The black contours identify data that were deemed significant by a permutation test (significance threshold $p<0.05$) and cluster corrected (cluster defining threshold $p<0.01$, see methods). The cluster-correction in two dimensions did not recover as significant all the same values along the diagonal that were deemed significant by the cluster-correction in one dimension in Figure 4b because it is more stringent. Nonetheless, the cross-temporal generalization for decoding hue showed greater significance further away from the diagonal compared to the cross-temporal generalization for luminance polarity (compare Figure 6a to Figure 6b). Figure 6c quantifies this comparison: yellow regions in the plot indicate where hue decoding was greater than luminance-polarity decoding, while dark blue regions indicate where luminance-polarity decoding was greater than hue decoding (the comparison was determined through a permutation test on the difference between luminance-polarity and hue decoding accuracy; red and white contours show cluster-corrected results). The relatively stronger cross-temporal generalization for hue, indicated in the plot as the yellow flanks around the diagonal, began relatively early after stimulus onset, emerging at about the same time as the initial peak in hue generalization decoding (119ms). The cross-temporal decoding plots also show evidence for faint bands of decoding parallel to the diagonal corresponding to a reactivation time of ~50ms, as described previously[23].

These results support the hypothesis that the patterns of activity in the brain associated with hue are more sustained than the patterns of activity associated with luminance polarity.

Luminance-polarity decoding with different combinations of hue for training and testing

The decoding analysis in Figure 3 shows that luminance polarity can be decoded from the pattern of MEG data but does not address the variability in the extent to which luminance-polarity information can be decoded using different combinations of hues for training and testing the classifiers. We were interested in addressing this question because the extent to which the brain relies on both L-M and S cone mechanisms for encoding luminance contrast is unclear[4, 5, 6], and because some behavioral data suggest that luminance is less reliably extracted from colors associated with the daylight locus (orange/blue) compared to the anti-daylight locus (pink/green)[65, 66, 67, 68, 69, 70]. For example, the chromaticity of natural lights is linked to relative

luminance contrast: direct sunlight has a warm spectrum (e.g. orange), and is associated with a luminance increment, whereas indirect light such as in shadows, has a cool spectrum (reflecting skylight) and is associated with a luminance decrement. If the visual system is adapted to natural statistics, we predicted that luminance polarity carried by orange and blue would be less meaningful about object boundaries than luminance polarity carried by green and pink, because luminance contrasts linked to orange and blue chromaticity are more likely to be attributed to the illumination.

Individual luminance-polarity decoding problems averaged over participants were variable both among the four identity problems (Figure 7a) and among the twelve generalization problems (Figure 7b). The color scale of the heatmap in Figure 7c indicates the luminance-polarity decoding performance for each of the 16 subproblems during the 5-ms time bin at the time of the initial peak decoding determined separately for each subproblem (these results are similar to those obtained during the 5 ms time bin set by the average peak decoding for the luminance-polarity identity problems). Data along the inverse diagonal correspond to the identity problems (recall the identity problems are those in which the hue was the same for training data and testing data). Data off the diagonal correspond to decoding problems for luminance polarity generalizing across hue, and these data allow us to test ideas about how S and L-M signals contribute to the representation of luminance polarity. The time bin of peak decoding for each problem is indicated in each entry. The bolded numbers are the total number of 5ms time bins in which decoding was significant (cluster-corrected, cluster-defining threshold $p<0.01$). Luminance polarity could be decoded above chance with data obtained in all combinations of training and test hues (significance of decoding was determined by a permutation test, see Methods), but there was variability in the magnitude and duration of decoding depending on the hue used to obtain data for training versus testing the classifiers (the 95% C.I. of peak decoding accuracy values obtained by bootstrapping are shown in square brackets).

The data in Figures 7 support three findings. First, luminance-polarity decoding was strongest when the carrier (training) hue was the same as the test hue (paired two-sided Wilcoxon signed-rank test on average peak decoding values comparing on-diagonal to off-diagonal Figure 7c, $p=0.0002$). Second, among these identity problems, classification performance was higher for warm colors (pink, orange) than cool colors (blue, green; repeated measures one-way ANOVA , Figure 7c, $p=0.003$; supporting our prior analysis[23]). Third, for a given pair of hues used to obtain the training and testing data, the performance of the classifiers was consistent regardless of which data set was used for training and which was used for testing;

this is reflected in a similar pattern in the upper and lower triangles of the heat map formed across the inverse diagonal in Figure 7c. For example, the accuracy for training on orange, testing on pink, is comparable to the accuracy for training on pink, testing on orange. But the magnitude of the classification accuracy and the number of significant time points recovered by the classifier varied with different combinations of hues for training/testing; for example, training (or testing) with orange and testing (or training) with pink yielded different results than training (or testing) with orange and testing (or training) with blue. We had hypothesized that classifiers trained and tested using hues associated with the daylight axis (blue and orange) would be less successful at decoding luminance polarity than classifiers trained and tested using hues associated with the anti-daylight axis (pink and green). The results show a trend supporting this hypothesis (Figure 7c), but it was not significant (repeated measures one-way ANOVA, Figure 7c, $p=0.076$).

Perhaps the appropriate test of our hypothesis should examine not the magnitude of decoding but rather the duration over which each problem could be decoded. The number of 5ms time bins that showed significant decoding was lower for problems involving hues of the daylight axis (34 or 36 time bins, bolded numbers, Figure 7c) compared to problems involving hues of the anti-daylight axis (42 or 45 bins), but this difference was not significant (Figure 7d shows the 95% CI of these values). Instead, an analysis of the duration of significant decoding shows that classifiers trained and tested with hues that differ in both their L-M and S cone modulation were worse than classifiers trained and tested with hues that differ in modulation of only one cone system (L-M or S). Figure 7d shows the results on the cone-opponent axes that define the color space, in which the color of the letters and arrows corresponds to the number of significant time bins for each problem (P, B, G, and O correspond to the four hues of the spirals: pink, blue, green, and orange; each problem is labeled with number of significant time bins and the 95% CI obtained by bootstrapping). The direction of the arrows indicates which data was used to train and test the classifiers; for example, the arrow from P to B shows the number of time bins in which the luminance polarity problems were significant when training on data obtained using light and dark pink, and tested on data obtained using light and dark blue. The arrows connecting hues through the color space, in which the hues for training and testing differ in both L-M and S modulation, showed a lower number of significant time bins compared to problems around the color circle, in which the hues for training and testing differ in only one or the other cone-opponent mechanism (the arrows across the diagonals are bluish, while all the problems around the perimeter are pink or gray). Thus the results do not provide strong support for our

439 initial hypothesis that neural mechanisms encoding luminance polarity are strongly impacted when the
440 hues of the stimuli correspond to the daylight locus, but they do support the idea that luminance-polarity is
441 encoded by a combination of both L-M and S chromatic mechanisms.
442
443 The statistical analysis shown in Figure 7 was obtained by comparing and averaging results across
444 participants. But the time course of decoding accuracy may show individual differences that will not be
445 apparent in these averages. One possibility we considered is whether the somewhat weaker luminance-
446 polarity decoding for blue and orange reflects variability in the timing and amplitude of decoding across
447 participants. To address this possibility, we first determined the average decoding accuracy throughout the
448 decoding time course (0-600ms) for each problem, for each participant (72 problems total: 18 participants x
449 2 daylight-axis problems x 2 anti-daylight-axis problems). The average peak decoding accuracy after
450 stimulus onset was not different for daylight colors and anti-daylight colors (55% [53.8, 57.0] versus 56%
451 [54.1, 57.8], p=0.68, repeated measures ANOVA). Next, for each problem, we calculated the Spearman
452 correlation of the decoding time course after stimulus onset for each of the 18 participants with the
453 average time course of a sample of 17 subjects drawn with replacement from the other participants, to
454 estimate the temporal correlation of the decoding curves. We repeated this 1000 times and averaged
455 across the repetitions to obtain a mean estimate of the temporal correlation of each participant's time
456 course to those of the others. If the time course of decoding shows the same structure across participants,
457 the analysis will yield high temporal correlations. The average correlation for problems across the daylight
458 colors versus across anti-daylight colors was not different (0.33 versus 0.30; p=0.62 repeated measures
459 ANOVA). Problems involving daylight colors were not different from problems involving anti-daylight colors
460 in both their mean accuracy and temporal correlation (MANOVA, p=0.37; note that the correlation of the
461 average decoding accuracy with the temporal correlation, Supplementary Figure 6, is not surprising,
462 because when decoding is higher, the temporal correlation will be more apparent).
463
464 *fMRI-guided MEG source localization*

465 We were interested in evaluating the extent to which MEG signals arising from functionally defined regions
466 in the cortex could support decoding of hue and luminance polarity. To address this question, we ran fMRI
467 experiments in 14 of the same participants in whom we collected MEG data, and we performed the MEG
468 analyses using subject-specific source localization. Our goal was to use functional data to define regions of

interest in the Ventral Visual Pathway (VVP) in each participant, controlling for individual differences in the absolute location of functional domains across people. In each subject we used fMRI to identify regions biased for faces, places, colors, and objects, using the same paradigm we used previously in which we measured fMRI responses to short movie clips of faces, bodies, objects, and scenes[71,72]. The paradigm involved measuring responses to intact and scrambled versions of the clips, and to clips in full color and black-and-white. As described in Lafer-Sousa et al (2016), the results can be used to define regions of interest: face-biased regions (including the FFA); place-biased regions (including the PPA); and color-biased regions. The results also recover area LO (Lateral Object area), defined by stronger responses to intact versus scrambled movie clips.

Figure 8a shows the fMRI results for one participant: greater responses to colored movie clips compared to black-and-white versions of the movie clips is shown by the heat map, apparent on the ventral view (Figure 8a, right panel); functional domains for faces (faces>objects), objects (intact objects>scrambled objects), and places (places>objects) are indicated by contour maps drawn at p=0.001 threshold. Color-biased activity was found sandwiched between place-biased activity (medially) and face-biased activity (laterally), confirming prior observations[72]. By aligning each participant to a standard atlas[73, 74], we also generated regions of interest for V1, V2, and MT, and for frontal cortex and the precentral gyrus (control regions).

MEG signals source localized to V1 and V2 yielded the highest magnitude modulation in current source density averaged across all stimulus presentations (Figure 8b). The magnitude of the CSD was different among the functional regions identified in the Ventral Visual Pathway (p=0.002, repeated measures one-way ANOVA): the color-biased regions were not different from face-biased regions (p=0.12; paired two sided t-test, uncorrected); but were different from LO (p=0.01), and from place-biased regions (p=0.02) (Figure 8c). These results provide a direct measure of neural activity, and confirm the indirect measurements obtained with fMRI suggesting that fMRI-identified color-biased regions (and possibly face-biased regions) play an important role in color processing[75].

Luminance polarity generalized across hue was decodable to some extent in all visual regions except the face-biased regions and the color-biased regions; it was most decodable in V1 and V2; to a lesser extent in

MT and LO; and to an even lesser extent in the place-biased regions; it was not decodable in the two control regions (Figure 8d). Hue generalized across luminance polarity was not reliably decodable in any region, but with a hint of decoding in V1 and/or V2 (Figure 8e). The distribution of sensors used in the decoding analysis is shown in Figure 8f, g (see legend and methods for details).

Discussion

The experiments here produced five results: first, hue and luminance polarity could be decoded independently. Luminance polarity could be decoded even if the training and testing stimuli differed in hue; and hue could be decoded even if the training and testing stimuli differed in luminance polarity. Second, classification accuracy was higher for problems in which the stimuli used to train and test the classifiers were identical (identity problems, Figure 3), compared to problems that required generalization across hue or luminance polarity (generalizing problems, Figure 3). Third, the time course of decoding was different for hue compared to luminance polarity (Figure 4). The peak for decoding hue was ~20 ms after the peak for decoding luminance polarity. Control experiments showed that these timing differences could not be attributed to the task that participants performed during MEG data collection (Figure 5). Fourth, representations of hue showed subtly greater cross-temporal generalization than representations of luminance polarity (Figure 6). And fifth, representations of luminance polarity varied depending on the hues used to obtain the training and testing data (Figure 7). Decoding of luminance polarity was most reduced when the hues used to obtain the training and testing data varied in both L-M and S, suggesting that both cone-opponent retinal mechanisms contribute to luminance encoding. Together, the results have implications for our understanding of the different roles that luminance polarity and hue play in visual perception, and how encoding these features is implemented by the neural circuitry.

Multivariate analyses of MEG and EEG data uncover important information about color processing in the brain[23, 76, 77, 78, 79, 80, 81]. The success of classifiers in decoding color implies that there is a spatial representation of color that gives rise to distinct spatial activation patterns in response to different colors. The success of the classifiers might be surprising given EEG and MEG are thought to show coarse spatial resolution, and the spatial scale of color columns, in early visual cortex [82, 83, 84, 85] or in higher visual cortex[86], is relatively fine. The paradox could be resolved either if the spatial scale of color columns is coarser than

527 currently estimated, or more likely, if the multivariate analyses are able to detect the existence of rather
528 fine spatial patterns, even if the analyses cannot recover the precise spatial organization of them[23, 87].
529
530 The present evidence for significant decoding of hue generalizing across luminance polarity, and luminance
531 polarity generalizing across hue, first presented in preliminary form[76, 77, 88], is consistent with results of a
532 recent independent report[89]. This evidence shows that either luminance contrast and hue are encoded to
533 some extent by different neural populations, or if they are encoded by entirely the same population, the
534 encoding must involve temporal multiplexing. Either possibility argues against the notion that luminance
535 polarity and hue are necessarily multiplexed simultaneously by the same neural population. The relatively
536 earlier timing of the luminance-polarity representation is consistent with univariate analyses of
537 electroencephalography data showing that the early onset P1 component is driven by luminance, and the
538 subsequent N1 peak is delayed by about 20 ms when elicited by equiluminant stimuli[39, 90, 91].
539
540 The difference in the timing of decoding hue versus luminance polarity was not impacted by task. The lack
541 of an impact of task is consistent with prior work showing that attentional tasks involving sequential stimuli
542 do not impact the feedforward flow of information[92], and task effects typically emerge 200-300 ms after
543 stimulus onset[62, 63]. Consistent with the late arrival of task effects, at times >300ms we observed subtly
544 higher decoding of luminance polarity compared to hue for data obtained with the 1-back luminance task;
545 and subtly higher decoding of hue compared to luminance polarity for data obtained with the 1-back hue
546 task (Figure 5c).
547
548 In object perception, the decoding time course reflects the perceptual and categorical dissimilarity of
549 stimuli, with more perceptually dissimilar stimuli (i.e. higher levels of abstraction) decodable later, and
550 associated with computations performed by areas further along the visual-processing hierarchy[59, 93, 94]. (An
551 increase in category abstraction is Doberman→dog→animal→animate.) One way of thinking about the
552 relatively later decoding of hue, then, is that (1) hue discrimination involves greater perceptual dissimilarity
553 (greater category abstraction) compared to luminance polarity; and (2) it is computed either by circuits
554 downstream of those that compute luminance polarity or requires more recurrent processing. The time
555 course for decoding hue (peak 118-119 ms; Figure 3b) is comparable to that for decoding shape-
556 independent object category[95] and face identity[96], operations that probably reflect activity in area LO and
557 the fusiform face area (FFA). The decoding time course for hue is therefore consistent with the hypothesis

558 that hue is computed at about the same distance from the retina along the visual processing hierarchy as
559 LO and the FFA, which implicates the posterior color-biased region of the Ventral Visual Pathway, part of
560 the V4 Complex[72, 97, 98, 99]. Cells in this region are spatially organized according to hue and their hue
561 selectivity is tolerant to changes in luminance polarity[36, 100, 101]. Another way of thinking about the results is
562 that the relatively later decoding of hue depends on a more sluggish channel from the retina through the
563 lateral geniculate nucleus up to visual cortex. This hypothesis is consistent with the idea that parvocellular
564 neurons underlie the color decoding, while magnocellular neurons underlie the luminance-polarity
565 decoding. Measurements of decoding in patients with cortical lesions that impact color perception may
566 prove instructive in addressing these alternative explanations[102].
567
568 Neural representations related to object vision that emerge earliest, as determined by classifiers tested
569 using the same images on which they were trained (i.e. not requiring perceptual dissimilarity), reflect the
570 encoding of "low-level visual features"—decoding in these classifiers peaks early (<100ms), and is
571 attributed to operations implemented early in the visual-processing hierarchy, perhaps V1[95, 103, 104, 105]. But
572 it has not been clear what constitutes the low-level visual features that determine the early decoding
573 performance. It is often implied that the low-level features consist of both oriented luminance-contrast
574 edges and color, as implied by separate luminance-contrast edge filters and color filters in the earliest layer
575 of convolutional neural networks[11]. But the timing discovered here suggests that (1) luminance contrast
576 (and not hue) is the first feature to be encoded; and (2) luminance contrast is implemented at an earlier
577 stage in the visual-processing hierarchy, before hue. The relative timing difference for decoding hue and
578 luminance polarity may be relatively small—just tens of milliseconds—but it is not only robust, evident
579 within individual subjects (Figure 5), but also long enough to encompass 7-10 synapses given monosynaptic
580 transmission latencies of ~2.5 ms in the visual pathway[106].
581
582 Knowledge of the timing of the neural events may provide clues to the different roles that luminance
583 contrast and hue play in vision. The visual system is confronted with a constant stream of retinal images,
584 each of which is associated with a cascade of neural activity lasting hundreds of milliseconds[60]. How does
585 the visual system parse this stream of information to encode new content? It must do so while retaining
586 some representations for longer durations to enable recognition and memory. In any situation in which
587 information is encoded in time, it is advantageous to have clear signals indicating the start and end of the
588 code, such as in genetics where gene sequences are parsed by canonical start and stop codons. In object

589 vision, dynamical systems modeling predicts the existence of observable update mechanisms that signal
590 new content[107]. The luminance-polarity decoding time course showed clear peaks corresponding to both
591 the onset and cessation of the stimulus (Figures 3, 4), and limited cross-temporal generalization (Figure 6a).
592 The time course for decoding hue, meanwhile, showed a weaker secondary peak and relatively stronger
593 cross-temporal generalization (Figure 6b, 6c). This pattern of results is consistent with the idea that the
594 brain more likely uses luminance contrast, rather than hue, as the updating signal, to encode discrete
595 events embedded in the constant stream of visual information.
596
597 What, then, does the timing for decoding hue suggest about its role in vision? People with normal color
598 vision find it difficult to make heterochromatic brightness matches[24] (see Supplementary Figure 5).
599 Moreover, humans are more likely to retain visual memory for hue than for luminance contrast[16]. Thus hue
600 is a more salient dimension than luminance contrast, and more sustained. The present results provide a
601 plausible neural correlate for these behaviors. They may also help resolve a paradox in visual
602 neurophysiology. In the earliest studies of V1 single-unit responses, Hubel and Wiesel remarked that "it
603 was surprising to us, however, that the great majority of cells [in V1] could discriminate precisely the
604 orientation or direction of movement of a stimulus, but had no marked selectivity regarding wave-
605 length"[108]. The seemingly weaker response of V1 cells to color compared to luminance contrast has fueled
606 a lot of research since this early observation[109]. The present results open the possibility that hue is encoded
607 by more sustained patterns of response to color across the population, which would not necessarily be
608 evident using traditional single-cell recording methods.
609
610 Decoding of luminance polarity varied depending on which hue was carrying the luminance-polarity signal
611 for training versus testing the classifiers. The experiments were designed to test our hypothesis that
612 luminance-contrast polarity is less readily encoded by the brain when carried by hues that align with the
613 daylight axis (orange/blue) versus when carried by colors that do not (pink/green)[67, 68, 69, 110]. The
614 chromaticity of natural lights covaries with relative luminance contrast, ranging from warm and bright (e.g.
615 direct sunlight) to cool and dark (e.g. shadow). Luminance contrast signals associated with blue and/or
616 orange chromaticities could be less reliable indicators of object boundaries since they may derive from
617 boundaries introduced by cast shadows. Consistent with this theory, the choices people make for what
618 constitutes "achromatic" vary most substantially along the daylight locus[65, 66]. The results provide partial
619 support for the hypothesis. Across participants, classifiers of luminance polarity trained and tested using

data obtained using blue performed the worst, and the difference in accuracy using classifiers trained and tested on blue compared to those trained and tested with orange was striking, which may relate to asymmetries in perception of blue-yellow color perception[68]. Classifiers trained and tested across daylight colors (blue and orange) trended towards lower accuracy and a shorter duration of significant decoding than classifiers trained and tested across anti-daylight colors (pink and green), but these differences were not significant. The MEG results therefore do not provide clear support for our initial hypothesis that the neural representation of luminance contrast is adapted to natural lighting conditions. Instead, the results show that decoding luminance polarity is most reduced when the hues for training and testing vary in both L-M and S (Figure 7d), suggesting that the brain encodes luminance polarity with both S and L-M mechanisms, which informs a long-standing discussion about which retinal mechanisms support the perception of luminance contrast[4, 5, 6].

There are substantial gaps in knowledge regarding the connection between fMRI and neural events. We leveraged the high spatial resolution of fMRI and the more direct access to neural events of MEG by obtaining both fMRI and MEG in the same participants. We used source localization to estimate MEG signals arising from fMRI-identified regions defined in each individual subject (Figure 8). The results provide a way of independently testing conclusions drawn from fMRI experiments. For example, within the Ventral Visual Pathway, cortical regions showing the strongest fMRI responses to color are sandwiched between more lateral regions responding most strongly to faces and more medial regions responding most strongly to places;[72] this pattern was described earlier in macaque monkeys[97]. Among regions of the VVP, the MEG signals assigned to color-biased regions showed the largest current-source density in response to the stimuli used in the MEG experiments—these stimuli differed only in color and not shape. Thus the present results support the idea that the VVP comprises parallel streams characterized by differential sensitivity to color information[38]. Source-localized analyses did not recover significant luminance-polarity-invariant hue representations, which may not be surprising given fundamental limitations of source localization[111]. In future work, we intend to use search-light analyses to probe for hue and luminance-polarity representations.

In summary, the present work used multivariate analyses of MEG data to quantitatively uncover the independent representations of hue and luminance polarity and their timing, which complements our

650 previous analyses of the combined, interacting, representations of hue and luminance polarity that underlie
651 the geometry of the neural representation of color[23]. The work reveals complex temporal dynamics that
652 mediate the representation of color, which presumably reflect the involvement of color in many visual
653 computations.
654
655 **Methods**
656 The MEGco data set (MEG colors data set) is published[23].
657
658 *Visual Stimuli*
659 Stimuli were eight square-wave spiral gratings on a neutral gray background (Figure 1a) [55, 56, 57]. The 8
660 stimulus colors, four hues at two luminance-contrast levels, of matched absolute cone contrast, were
661 defined in DKL color space[16, 52, 53] using implementations by Westland and Brainard : the axes of this color
662 space are defined in terms of activation of the two cone-opponent post-receptoral chromatic mechanisms
663 (Supplementary Figure 1). The z-axis is defined by luminance contrast. The four hues were defined by the
664 intermediate axes of DKL space: at 45° (pink), 135° (blue), 225° (green), and 315° (orange). The absolute L-
665 M cone modulation of the stimuli were matched; and the absolute S cone modulation of the stimuli were
666 matched. Different colors were created by pairing different signs of modulation along L-M and S. Two
667 spirals – one positive luminance contrast (+20° elevation; "light") and one negative luminance contrast (-
668 20° elevation; "dark") – were created at each hue. The neutral adapting background was 33.5 cd/m2. The
669 luminance contrast of the stimuli was 25%, computed as Weber contrast because the stimuli were brief and
670 not full field. Modulation of the cone-opponent mechanisms, shown in Supplementary Figure 1, was
671 computed relative to the adapting background gray. Using MATLAB scripts of Westland et al. , the xyY
672 values of the stimuli were converted to LMS values using *xy2MB*, and the LMS values were converted to
673 DKL values using *lms2dkl*.
674
675 We designed the experiment such that (i) the stimuli were well above detection threshold, and (ii) the
676 absolute cone contrast of different hues (e.g. light pink versus light green) was comparable to the absolute
677 cone contrast of different luminance contrast levels (e.g. light pink versus dark pink). We can estimate the
678 contrast of the stimuli in units of detection threshold, using detection data provided by Sachtler and Zaidi
679 (1992)[16]. It is important to do this separately for color and luminance contrast because, as Sachtler and
680 Zaidi showed, detection thresholds (measured in absolute cone contrast) are over 9x higher for luminance

contrast than for color. The cone contrast is computed as the absolute value of the contrast in L, plus the absolute value of the contrast in M. For example:

We compute the luminance contrast carried by the pink hue as follows:

(equation 1)

$$|(Lcone_{light_pink} - Lcone_{dark_pink})/(Lcone_{light_pink} + Lcone_{dark_pink})| + |(Mcone_{light_pink} - Mcone_{dark_pink})/(Mcone_{light_pink} + Mcone_{dark_pink})|$$

$$= |(27.6333-16.5006)/(27.6333+16.5006)| + |(14.2617-8.1970)/(14.2617+8.1970)|$$

$$= 0.5225$$

The detection threshold (|L|+|M|) for luminance contrast is 0.014 (from Sachtler and Zaidi[16]). Thus the contrast of the stimuli was (0.5225/0.014), ~37x detection threshold.

We can compute the color contrast between pink and green at a given luminance contrast as follows:

(equation 2)

$$|(Lcone_{light_pink} - Lcone_{light_green})/(Lcone_{light_pink} + Lcone_{light_green})| + |(Mcone_{light_pink} - Mcone_{light_green})/(Mcone_{light_pink} + Mcone_{light_green})|$$

$$= 0.0525$$

The detection threshold (|L|+|M|) for color contrast is 0.001724 (from Sachtler and Zaidi[16]). Thus the contrast of the light pink and green stimuli we used is (0.0525/0.001724), ~30 x detection threshold. Comparable calculations for the dark pink and dark green stimuli show these stimuli to be ~51X detection threshold. The two values of color contrast (for light stimuli and dark stimuli) straddle the absolute contrast (in detection units) across luminance (37x, as calculated above). The possible difference in color contrast of light versus dark stimuli had no impact on the time course for decoding hue (see Supplementary Figure 4),

711 showing that at the supra-threshold values used, decoding strength is not likely influenced by stimulus
712 saturation.
713
714 *MEG Acquisition and Preprocessing*
715
716 Participants were scanned in the Athinoula A. Martinos Imaging Center of the McGovern Institute for Brain
717 Research at the Massachusetts Institute of Technology (MIT) over the course of 2 sessions, on an Elekta
718 Triux system (306-channel probe unit consisting of 102 sensor triplets, with 204 planar gradiometer
719 sensors, and 102 magnetometer sensors). The experimental paradigm was created using Psychtoolbox[114];
720 stimuli were back-projected onto a 44" screen using a SXGA+ 10000 Panasonic DLP Projector, Model No.
721 PT-D10000U (50/60Hz, 120V). Data was recorded at a sampling rate of 1000Hz, filtered between 0.03-
722 330Hz. Head location was recorded by means of 5 head position indicator (HPI) coils placed across the
723 forehead and behind the ears. Before the MEG experiment began, 3 anatomical landmarks (bilateral
724 preauricular points and the nasion) were registered with respect to the HPI coils, using a 3D digitizer
725 (Fastrak, Polhemus, Colchester, Vermont, USA). During recording, pupil diameter and eye position data
726 were collected simultaneously using an Eyelink 1000 Plus eye tracker (SR Research, Ontario, Canada) with
727 fiber optic camera.
728
729 Once collected, raw data was preprocessed to offset head movements and reduce noise by means of
730 spatiotemporal filters (Taulu et al, 2004; Taulu & Simola, 2006), with Maxfilter software (Elekta,
731 Stockholm). Default parameters were used: harmonic expansion origin in head frame = [0 0 40] mm;
732 expansion limit for internal multipole base = 8; expansion limit for external multipole base = 3; bad
733 channels omitted from harmonic expansions = 7 s.d. above average; temporal correlation limit = 0.98;
734 buffer length = 10 s). In this process, a spatial filter was applied to separate the signal data from noise
735 sources occurring outside the helmet, then a temporal filter was applied to exclude any signal data highly
736 correlated with noise data over time. Following this, Brainstorm software (Tadel et al., 2011) was used to
737 extract the peri-stimulus MEG data for each trial (-200 to 600 ms around stimulus onset) and to remove the
738 baseline mean.
739
740
741 *MEG Participants and Task*

All participants (N=18, 11 female, age 19-37 years) had normal or corrected-to-normal vision, were right-handed, spoke English as a first language, and provided informed consent. One participant was an author and thus not naïve to the purpose of the study. During participants' first session, they were screened for colorblindness using Ishihara plates; they also completed a version of a color-naming task as part of a separate study. After this task, participants completed a 100-trial practice session of the 1-back hue-matching task that would be used in the MEG experimental sessions. Once this was complete, participants were asked if they had any questions about the task or the experiment; eye-tracking calibration was performed; and MEG data collection began.

During the MEG data collection, participants were instructed to fixate at the center of the screen. Spirals were presented subtending 10° of visual angle for 116 ms, centered on the fixation point. The fixation point was a white circle that appeared during inter-trial intervals (ITIs, 1s). In addition to the spirals, the words "green" and "blue" were presented in white on the screen for the same duration, and probe trials were presented with a white "?". (Responses to the words were analyzed as part of a separate study.) During the probe trials, which occurred every 3-5 stimulus trials (pseudo-randomly interspersed, 24 per run), participants were instructed to report via button press if the two preceding spirals did or did not match according to hue (1-back hue task). Maximum response time was 1.8s, but the trials advanced as soon as participants answered.

Participants were encouraged to blink only during probe trials, as blinking generates large electrical artifacts picked up by the MEG. Each run comprised 100 stimulus presentations, and participants completed 25 runs per session over the course of approximately 1.5 hours. Between each run, participants were given a break to rest their eyes and speak with the researcher if necessary. Once 10s had elapsed, participants chose freely when to end their break by button-press. Over the course of both sessions, participants viewed each stimulus 500 times. Individual runs were identical across subjects, but the order of runs was randomized between subjects. The sequence of stimuli within each run was random with the constraint that the total number of presentations was the same for each stimulus condition over the set of runs obtained for each participant.

In the main experiments, data from all participants was used (no data was excluded because of poor

773 behavioral performance).

774

775 All experimental procedures involving human participants were approved by the Wellesley College
776 Institutional Review Board, the Massachusetts Institute of Technology Committee on the Use of Humans as
777 Experimental Subjects, and the National Institutes of Health Intramural Institute Clinical Research Review
778 Committee. Participants consented to participate and to having their data published.

779

780 *Pilot experiment to determine the number of trials per condition*

781

782 Prior to conducing the main experiments, we collected data in four participants (3 female), using four
783 colored stimuli (two hues, blue and orange, at light and dark contrast levels). Participants were instructed
784 to fixate a small spot at the center of the display; besides this passive fixation, there was no task. Eye
785 movements were monitored to ensure passive fixation. The goal of this pilot was to determine the number
786 of trials needed to successfully decode color. As in the main experiment, the stimuli were spirals, subjects
787 viewed each color 500 times, and each stimulus appeared for 100 ms followed by a 1 s ISI. Other details of
788 experimental paradigm were the same as for the main experiment, except that the stimuli included only
789 two hues at two luminance polarity values (four conditions total). The two hues were DKL angles 150 and
790 300, which are intermediate colors corresponding roughly to blue and yellow; the luminance of the stimuli
791 were: background gray, 41 cd/m^2; positive luminance-contrast stimuli, 48-50 cd/m^2; and negative
792 luminance-contrast stimuli, 30-32 cd/m^2). To evaluate the experimental power, we set out to present each
793 stimulus with what we thought was an excessive number of trials: 500 (leaving 375 after rejecting trials
794 with eye blinks or other artifacts). Supplementary Figure 2 shows the data reliability for the pilot
795 experiment. The figure was generated by subsampling independent pairs of N% of data, computing the
796 decoding for each independent set of data, and computing the correlation coefficient between the sets of
797 data at each time point in the decoding curve. As expected, the test-retest curves plateau; but the plateau
798 occurs only with substantial numbers of trials. Informed by these pilot experiments, we performed the
799 main experiment with 500 trials. Preliminary results of these experiments have been presented and provide
800 to our knowledge the first evidence for decoding color from MEG data[76, 77]. The four participants who took
801 part in this pilot experiment did not participate in the main experiments or the control experiment.

802

803 *Control experiment to measure the impact of task*

Before launching the main experiments, we deployed a control experiment to determine the decoding parameters for the main experiments and to evaluate the task effects on decoding. This experiment was conducted in 2 participants (1 female, age 20-30 years) who completed 5 sessions of 20 runs each, one of which had a technical glitch and was discarded. Each run had 100 stimulus presentations. Every 3-5 stimulus presentations, participants saw a "?" that prompted a response regarding the preceding two stimuli; the correct answer depended on the task for the run. For one half of the runs, the participants performed a 1-back hue matching task, and during the other half of the runs they performed a 1-back luminance-polarity matching task (Figure 5a). Participants performed the same task for trials in a run, with the task alternating between runs for a given session. The results from the control experiment showed no impact of task performance on the latency or time-to-peak of decoding (Figure 5 b,c). For consistency in the main experiment, we had participants perform the same task. The decoding analysis shown in Figure 5b,c, is for the identity problems illustrated in Figure 1. The two participants who took part in the control experiment did not participate in the main experiments or the other pilot experiment. Exact p values are reported throughout, except in for tests based on bootstrapping, where the number of bootstrap comparisons is zero. In those cases, the p value can only be said to be less than 0.001 because there were 1000 bootstraps.

MEG Processing and Decoding Analyses

Brainstorm software was used to process MEG data. Trials were discarded if they contained eyeblink artifacts, or contained out-of-range activity in any of the sensors (0.1-8000 fT). Three participants exhibited sensor activity consistently out of range, so this metric was not applied to their data as it was not a good marker of abnormal trials. After excluding bad trials, there were at least 375 good trials for every stimulus type for every participant. Data were subsampled as needed to ensure the same number of trials per condition were used in the analysis.

Decoding was performed using the Neural Decoding Toolbox (NDT)[54]. We used the maximum correlation coefficient classifier in the NDT to train classifiers to associate patterns of MEG activity across the sensors with the visual stimuli presented. This classifier computes the mean population vector for sets of trials belonging to each class in the training data and calculates the Pearson's correlation coefficient between

those vectors and the test vectors. The class with the highest correlation is the classifier's prediction. The main conclusions were replicated when using linear support vector machine classifiers. The classifiers were tested using held-out data—i.e. data that was not used in training. Data from both magnetometers and gradiometers were used in the analysis, and data for each sensor was averaged into 5-ms non-overlapping bins from 200 ms before stimulus onset to 600 ms after stimulus onset.

Custom MATLAB code was used to format MEG data preprocessed in Brainstorm for use in the NDT and to combine the two data-collection sessions for each participant. Decoding was performed independently for each participant, and at each time point. As illustrated in Figure 2, for each decoding problem, at each timepoint (a 5 ms time bin), the 375 trials for each stimulus condition were divided into 5 sets of 75 trials. Within each set, the 75 trials were averaged together. This process generated 5 cross-validation splits: the classifier was trained on four of these sets, and tested on one of them, and the procedure was repeated five times so that each set was the test set once. This entire procedure was repeated 50 times, and decoding accuracies reported are the average accuracies across these 50 decoding "runs". This procedure ensured that the same data was never used for both training and testing, and it also ensured the same number of trials was used for every decoding problem. The details of the cross-validation procedure, such as the number of cross-validation splits, were determined during the pilot experiments to be those that yielded a high signal-to-noise ratio (SNR) and high decoding accuracy in both participants on the stimulus identity problem.

On each run, both the training and test data were z-scored using the mean and standard deviation over all time of the training data. Following others, we adopted a de-noising method that involved selecting for analysis data from the most informative sensors[115]; we chose the 25 sensors in the training data whose activity co-varied most significantly with the training labels. These sensors were identified as those with the lowest p-values from an F-test generated through an analysis of variance (ANOVA); the same sensors were then used for both training and testing. The sensor selection was specific for each participant. The sensors chosen tended to be at the back of the head (Figure 8f,g). Analyses using all channels, rather than selecting only 25, yielded similar results.

All classification problems were binary (see Figures 1 and 2). For each problem, a classifier was trained and tested in 5ms bins from time t=200ms before stimulus onset to t=600ms after stimulus onset (Figure 2).

The classifiers' performance shown in Figures 3 and 4 were generated through a bootstrapping procedure. First, the problems were evaluated for each participant, resulting in 18 independent decoding time courses for each unique problem. The decoding time courses for each problem were sampled 18 times with replacement and averaged, and this procedure was repeated 1000 times to produce 1000 time courses, which were averaged to generate the decoding traces in Figures 3 and 4.

The 95% CI around the time to peak was determined from the times to peak for each of the 1000 bootstrapped time courses. Statistical tests on the difference in time to peak between two problems were performed using the bootstrap distributions of the differences in time to peak values. If the 95% confidence interval did not include 0, we rejected the null hypothesis, and p values were calculated based on the proportion of values that did fall below 0. To determine at which time points decoding accuracy was significantly above chance, a permutation test was used to calculate p values (Pantazis et al, 2005). This was done by permuting the sign of the decoding accuracy data on a participant basis 1000 times. For each permutation sample, the mean accuracy was recomputed, resulting in an empirical distribution of 1000 mean accuracies. This distribution was used to convert the real mean accuracies across subjects over time to p-value maps over time. The significance threshold was $p<0.05$, and significant regions were determined using a cluster-based approach (cluster defining threshold $p<0.01$)[59, 116]. Onset of significance was calculated as the first time point where accuracy was significant for four continuous 5-ms time bins—the requirement that the accuracy be significant for four consecutive bins was adopted to minimize false positives. Reported p values for paired comparisons of the timing and magnitude of the decoding accuracy are uncorrected.

For the control experiment, classifiers were trained and tested for each session individually to yield eight decoding time courses; the 5 cross validation splits consisted of 15 trials. The 95% confidence intervals on the times to peak and the statistical tests used to compare the peak times were produced similarly to those of the main experiment, by sampling from the 8 decoding curves to produce 1000 bootstrapped time courses (shading in Figure 5 shows the 95% CI at each time point). To estimate the significance of decoding at each time point, shown as the horizontal line of data points above the x-axis in the curves in Figure 5, we used a within-subjects approach instead of the across subjects approach used in the main experiment. A null distribution was produced to test whether the decoding at each time point was greater than what one would expect by chance. In each null decoding run, the stimuli identities were shuffled, and the classifier

897 was trained and tested on the shuffled data. There were 1000 null decoding runs, and p values at each time
898 point were calculated based on the proportion of null decoding accuracies that exceeded the real decoding
899 accuracy. These p values were then FDR corrected.
900
901 The test-retest curves to evaluate experimental power (Supplementary Figures 2 and 3) were obtained by
902 drawing pairs of independent samples of 10%, 25%, 40%, and 50% of the trials (from a total of 375 trials),
903 determining the correlation of the classification performance among the subproblems between the pairs,
904 and repeating the procedure 5 times to generate error bars. For example, for the "10%" data point in the
905 graph, we drew two sets of 10% of the trials at random—no trials were common to both sets. We trained
906 separate classifiers on each of the independent sets and computed the temporal correlation between the
907 two decoding time courses. We repeated this procedure 5x. We then averaged the correlation coefficients
908 from the 5 repeats to obtain error bars.
909
910 In Figure 6, we tested the performance of the classifiers across time: each classifier trained using data
911 obtained at each time bin was tested using data obtained at every 5 ms time bin from -200 to 600 ms after
912 stimulus onset creating a 2-dimensional matrix of decoding results. Significant time points were determined
913 using a sign-permutation test and cluster correction (cluster-defining threshold $p<0.01$)[59, 116]. In Figure 6c to
914 compare the decoding of hue and luminance, the sign-permutation was carried out on a subproblem basis.
915
916 In Figure 7, the 95%CI of the classification performance for each problem (the entries in Figure 7b, 7c) were
917 generated by bootstrapping across participants (n=1000).
918
919 Note that decoding using pupil diameter and eye position cannot account for decoding of color, as
920 discussed in a previous analysis of the MEGco data set[23].
921
922 *MRI dynamic Localizer Task*
923
924 To localize shape, place, face, and color-biased regions of interest (ROIs), 14 of 18 participants were
925 scanned using the fMRI dynamic Localizer (DyLoc) described in Lafer Sousa et al, 2016, with the same
926 parameters described there. In brief, participants passively viewed full color and grayscale (achromatic)
927 versions of natural video clips that depicted faces, bodies, scenes, objects, and scrambled objects.

928 Scrambled objects clips were clips in the object category that were divided into a 15 by 15 grid covering the
929 frame, the boxes of which were then scrambled. Participants completed 8 runs of the task, each of which
930 contained 25 blocks of 18 s (20 stimuli and 5 gray fixation blocks). The stimuli were a maximum of 20° of
931 visual angle wide and 15° tall. A Siemens 3T MAGNETOM Prisma fit scanner (Siemans AG, Healthcare,
932 Erlangen, Germany) with 64 RF receivers in the head coil was used to collect MRI data in 8 of 14
933 participants, while a Siemens 3T MAGENTOM Tim Trio scanner with 32 channels in the head coil was used
934 for the other 6 subjects.
935
936 For both groups, following Lafer-Sousa et al, a T2*-weighted echo planar imaging (EPI) pulse sequence was
937 used to detect blood-oxygen-level-dependent (BOLD) contrast. Field maps (2 mm isotropic, 25 slices) were
938 collected before each DyLoc run for the purpose of minimizing spatial distortions due to magnetic
939 inhomogeneities in the functional volumes during analysis. Functional volumes (2 mm isotropic, 25 slices,
940 field of view [FOV] = 192 mm, matrix = 96x96mm, 2.0 s TR, 30 ms TE, 90˚ flip angle, 6/8 echo fraction) were
941 collected on a localized section of the brain, aligned roughly parallel to the temporal lobe. The volumes
942 covered V1-V4 in occipital cortex as well as the entirety of the temporal lobe ventral to the superior
943 temporal sulcus (STS), and in some cases including parts of the STS. To allow for T1 equilibration, in each
944 run, the first 5 volumes were not used during analysis.
945
946 High-resolution T1-weighted anatomical images were also collected for each subject by means of a
947 multiecho MPRAGE pulse sequence (1 mm isotropic voxels, FOV = 256 mm, matrix = 256x256mm).
948
949 *MRI Analysis*
950
951 MRI data were processed following Lafer-Sousa et al (2016). Using Freesurfer
952 (https://surfer.nmr.mgh.harvard.edu) and custom MATLAB scripts, the anatomical volumes were segmented
953 into white- and gray-matter structures (Dale et al, 1999; Fischl et al, 1999, 2001). Functional data,
954 processed on an individual subject basis, were field- and motion-corrected (by means of rigid-body
955 transformations to the middle of each run), normalized for intensity after masking non-brain tissue, and
956 spatially smoothed with an isotropic Gaussian kernel (3 mm FWHM) for better SNR. Subsequently,
957 Freesurfer's bbregister was used to generate a rigid-body transformation used to align the functional data
958 to the anatomical volume.

Whole-volume general linear model-based analyses were performed for all 8 runs collected for each participant, using boxcar functions convolved with a gamma hemodynamic response function as regressors (Friston et al, 1994); each condition's boxcar function included all blocks from that condition, as well as nuisance regressors for motion (three translations, three rotations) and a linear trend to capture slow drifts.

Brain regions used to restrict decoding analyses of MEG source data were defined using two methods. Anatomically defined regions were defined using surface-based Freesurfer atlases: "precentral" and "frontal" regions corresponded respectively to the "precentral" and "rostralmiddlefrontal" bilateral regions in the Desikan-Killiany atlas (Desikan et al, 2006); V1 (BA 17), V2 (BA 18) and MT regions were defined using the Brodmann atlas (https://surfer.nmr.mgh.harvard.edu/fswiki/BrodmannAreaMaps ; Brodmann, 1909). Functionally defined regions were defined individually using Lafer-Sousa et al (2016) as a reference. FFA was selected from voxels where face response>object response ($p<0.001$), using data from all 8 runs. The same procedure was followed for VVP-c from voxels where color response>grayscale response, PPA from voxels where scene response>object response, and LO from voxels where object response>scrambled object response.

Source Localization and Decoding with ROIs

Current source density is a metric representing the current at each point on the surface of the brain, defined by the source grid. First, using Brainstorm, a minimum norm estimate (MNE) was calculated, which was "depth-weighted", to compensate for a bias in current density calculations that results in more activity being placed on superficial gyrii, neglecting regions of cortex embedded in deeper sulci (Hämäläinen, 2009). The MNE at a given source was normalized by the square root of a local estimate of noise variance (dynamical Statistical Parametric Mapping; dSPM, Dale, 2000), yielding a unitless z-scored statistical map of activity. Once a source map was created, ROI analysis was performed by restricting the features of the classifiers to the top 25 sources within the bounds of a given ROI whose activity covaried most with the training labels, using custom code. Additionally, the sources within an ROI were averaged together within subjects to yield the average sensor response by ROI.

990 *Data Availability*

991 Source data are provided with the paper. The MEGco dataset[23] can be accessed at the OpenNeuro data
992 base and the NEICOMMONS databases (https://doi.org/10.18112/openneuro.ds003352.v1.0.0;
993 https://neicommons.nei.nih.gov/#/).

994

995 *Code Availability*

996 The procedures and code used to produce all the figures, and all statistical analyses, in the current study
997 will be available in the NEICOMMONS repository.

998

999

1000 **References**

1001

1002 1. Kuehni R, Schwartz A. Color Ordered: A Survey of Color Systems from Antiquity to the Present. *New York, NY:*
1003 *Oxford University press*, (2008).

1004

1005 2. Smet KA, Webster MA, Whitehead LA. A simple principled approach for modeling and understanding uniform
1006 color metrics. *J Opt Soc Am A Opt Image Sci Vis* **33**, A319-331 (2016).

1007

1008 3. Ennis RJ, Zaidi Q. Geometrical structure of perceptual color space: Mental representations and adaptation
1009 invariance. *J Vis* **19**, 1 (2019).

1010

1011 4. Lee J, Stromeyer CF, 3rd. Contribution of human short-wave cones to luminance and motion detection. *J*
1012 *Physiol* **413**, 563-593 (1989).

1013

1014 5. Ripamonti C, Woo WL, Crowther E, Stockman A. The S-cone contribution to luminance depends on the M-
1015 and L-cone adaptation levels: silent surrounds? *J Vis* **9**, 10 11-16 (2009).

1016

1017 6. Conway BR. Color signals through dorsal and ventral visual pathways. *Vis Neurosci* **31**, 197-209 (2014).

1018

1019 7. Favreau OE, Cavanagh P. Color and luminance: independent frequency shifts. *Science* **212**, 831-832 (1981).

1020

1021 8. Krauskopf J, Williams DR, Heeley DW. Cardinal directions of color space. *Vision Research* **22**, 1123-1131
1022 (1982).

1023

1024 9. Goddard E, Chang DHF, Hess RF, Mullen KT. Color contrast adaptation: fMRI fails to predict behavioral
1025 adaptation. *NeuroImage* **201**, 116032 (2019).

1026

1027 10. Dobkins K. Moving colors in the lime light. *Neuron* **25**, 15-18 (2000).

1028
1029 11. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. *in*
1030 *Advances in Neural Information Processing systems*, 1097-1105 (2012).

1031
1032 12. Zeiler MD, Fergus R. Visualizing and Understanding Convolutional Networks. *Computer Vision - Eccv 2014, Pt I*
1033 **8689**, 818-833 (2014).

1034
1035 13. Rafegas I, Vanrell M. Color representation in CNNs: parallelisms with biological vision. *2017 Ieee International*
1036 *Conference on Computer Vision Workshops (Iccvw 2017)*, 2697-2705 (2017).

1037
1038 14. Flachot A, Gegenfurtner KR. Processing of chromatic information in a deep convolutional neural network. *J*
1039 *Opt Soc Am A Opt Image Sci Vis* **35**, B334-B346 (2018).

1040
1041 15. Hansen T, Gegenfurtner KR. Independence of color and luminance edges in natural scenes. *Vis Neurosci* **26**,
1042 35-49 (2009).

1043
1044 16. Sachtler WL, Zaidi Q. Chromatic and luminance signals in visual memory. *J Opt Soc Am A* **9**, 877-894 (1992).

1045
1046 17. Clifford CW, Spehar B, Solomon SG, Martin PR, Zaidi Q. Interactions between color and luminance in the
1047 perception of orientation. *J Vis* **3**, 106-115 (2003).

1048
1049 18. Yaguchi H, Ikeda M. Subadditivity and superadditivity in heterochromatic brightness matching. *Vision Res* **23**,
1050 1711-1718 (1983).

1051
1052 19. Buck SL. Brown. *Curr Biol* **25**, R536-537 (2015).

1053
1054 20. Lindsey DT, Brown AM. World Color Survey color naming reveals universal motifs and their within-language
1055 diversity. *Proc Natl Acad Sci U S A* **106**, 19785-19790 (2009).

1056
1057 21. Gibson E*, et al.* Color naming across languages reflects color use. *Proceedings of the National Academy of*
1058 *Sciences of the United States of America* **114**, 10785-10790 (2017).

1059
1060 22. Zaslavsky N, Kemp C, Tishby N, Regier T. Color Naming Reflects Both Perceptual Structure and
1061 Communicative Need. *Top Cogn Sci* **11**, 207-219 (2019).

1062
1063 23. Rosenthal IA, Singh SR, Hermann KL, Pantazis D, Conway BR. Color Space Geometry Uncovered with
1064 Magnetoencephalography. *Curr Biol*, (2020).

1065
1066 24. De Valois KK, Switkes E. Simultaneous masking interactions between chromatic and luminance gratings.
1067 *Journal of the Optical Society of America* **73**, 11-18 (1983).

1068
1069 25. Maunsell JH, Ghose GM, Assad JA, McAdams CJ, Boudreau CE, Noerager BD. Visual response latencies of
1070 magnocellular and parvocellular LGN neurons in macaque monkeys. *Vis Neurosci* **16**, 1-14 (1999).
1071
1072 26. Gegenfurtner KR. Cortical mechanisms of colour vision. *Nat Rev Neurosci* **4**, 563-572 (2003).
1073
1074 27. Thorell LG, De Valois RL, Albrecht DG. Spatial mapping of monkey V1 cells with pure color and luminance
1075 stimuli. *Vision Research* **24**, 751-769 (1984).
1076
1077 28. Nealey TA, Maunsell JH. Magnocellular and parvocellular contributions to the responses of neurons in
1078 macaque striate cortex. *J Neurosci* **14**, 2069-2079 (1994).
1079
1080 29. Johnson EN, Hawken MJ, Shapley R. Cone inputs in macaque primary visual cortex. *J Neurophysiol* **91**, 2501-
1081 2514 (2004).
1082
1083 30. Lennie P, Krauskopf J, Sclar G. Chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience*
1084 **10**, 649-669 (1990).
1085
1086 31. Conway BR. Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1).
1087 *Journal of Neuroscience* **21**, 2768-2783 (2001).
1088
1089 32. Horwitz GD, Hass CA. Nonlinear analysis of macaque V1 color tuning reveals cardinal directions for cortical
1090 color processing. *Nat Neurosci* **15**, 913-919 (2012).
1091
1092 33. Nassi JJ, Callaway EM. Parallel processing strategies of the primate visual system. *Nat Rev Neurosci* **10**, 360-
1093 372 (2009).
1094
1095 34. Sincich LC, Horton JC. Divided by cytochrome oxidase: a map of the projections from V1 to V2 in macaques.
1096 *Science* **295**, 1734-1737 (2002).
1097
1098 35. Verhoef BE, Bohon KS, Conway BR. Functional architecture for disparity in macaque inferior temporal cortex
1099 and its relationship to the architecture for faces, color, scenes, and visual field. *J Neurosci* **35**, 6952-6968
1100 (2015).
1101
1102 36. Bohon KS, Hermann KL, Hansen T, Conway BR. Representation of Perceptual Color Space in Macaque
1103 Posterior Inferior Temporal Cortex (the V4 Complex). *Eneuro* **3**, (2016).
1104
1105 37. Nasr S, Polimeni JR, Tootell RB. Interdigitated Color- and Disparity-Selective Columns within Human Visual
1106 Cortical Areas V2 and V3. *J Neurosci* **36**, 1841-1857 (2016).
1107
1108 38. Conway BR. The Organization and Operation of Inferior Temporal Cortex. *Annu Rev Vis Sci* **4**, 381-402 (2018).
1109

1110 39. Rabin J, Switkes E, Crognale M, Schneck ME, Adams AJ. Visual evoked potentials in three-dimensional color
1111 space: correlates of spatio-chromatic processing. *Vision Res* **34**, 2657-2671 (1994).

1112
1113 40. Kuriki I, Sadamoto K, Takeda T. MEG recording from the human ventro-occipital cortex in response to
1114 isoluminant color stimulation. *Vis Neurosci* **22**, 283-293 (2005).

1115
1116 41. Logothetis NK, Schiller PH, Charles ER, Hurlbert AC. Perceptual deficits and the activity of the color-opponent
1117 and broad-band pathways at isoluminance. *Science* **247**, 214-217 (1990).

1118
1119 42. Shevell SK, Kingdom FA. Color in complex scenes. *Annu Rev Psychol* **59**, 143-166 (2008).

1120
1121 43. Engel SA. Adaptation of Oriented and Unoriented Color-Selective Neurons in Human Visual Areas. *Neuron* **45**,
1122 613-623 (2005).

1123
1124 44. Wade A, Augath M, Logothetis N, Wandell B. fMRI measurements of color in macaque and human. *J Vis* **8**, 6
1125 1-19 (2008).

1126
1127 45. Mullen KT, Thompson B, Hess RF. Responses of the human visual cortex and LGN to achromatic and
1128 chromatic temporal modulations: an fMRI study. *J Vis* **10**, 13 (2010).

1129
1130 46. Boynton GM, Engel SA, Heeger DJ. Linear systems analysis of the fMRI signal. *NeuroImage* **62**, 975-984
1131 (2012).

1132
1133 47. Switkes E, Crognale MA. Comparison of color and luminance contrast: apples versus oranges? *Vision Res* **39**,
1134 1823-1831 (1999).

1135
1136 48. Hardman AC, Martinovic J. Salience of spatiochromatic patterns. *J Vis* **21**, 7 (2021).

1137
1138 49. Connah D, Finlayson GD, Bloj M. Seeing beyond luminance: A psychophysical comparison of techniques for
1139 converting colour images to greyscale. *Fifteenth Color Imaging Conference: Color Science and Engineering*
1140 *Systems, Technologies, and Applications, Final Program and Proceedings*, 336-+ (2007).

1141
1142 50. Kanan C, Cottrell GW. Color-to-grayscale: does the method matter in image recognition? *PLoS One* **7**, e29740
1143 (2012).

1144
1145 51. Hebart MN, Baker CI. Deconstructing multivariate decoding for the study of brain function. *NeuroImage* **180**,
1146 4-18 (2018).

1147
1148 52. MacLeod DI, Boynton RM. Chromaticity diagram showing cone excitation by stimuli of equal luminance. *J Opt*
1149 *Soc Am* **69**, 1183-1186 (1979).

1150

1151 53. Derrington AM, Krauskopf J, Lennie P. Chromatic mechanisms in lateral geniculate nucleus of macaque.
1152 *Journal of Physiology* **357**, 241-265 (1984).

1153
1154 54. Meyers EM. The neural decoding toolbox. *Front Neuroinform* **7**, 8 (2013).

1155
1156 55. Mannion DJ, McDonald JS, Clifford CW. Discrimination of the local orientation structure of spiral Glass
1157 patterns early in human visual cortex. *NeuroImage* **46**, 511-515 (2009).

1158
1159 56. Seymour K, Clifford CW, Logothetis NK, Bartels A. Coding and binding of color and form in visual cortex. *Cereb*
1160 *Cortex* **20**, 1946-1954 (2010).

1161
1162 57. Brouwer GJ, Heeger DJ. Categorical clustering of the neural representation of color. *J Neurosci* **33**, 15454-
1163 15465 (2013).

1164
1165 58. Norman-Haignere SV, Kanwisher N, McDermott JH, Conway BR. Divergence in the functional organization of
1166 human and macaque auditory cortex revealed by fMRI responses to harmonic tones. *Nature Neuroscience*
1167 **22**, 1057-+ (2019).

1168
1169 59. Cichy RM, Pantazis D, Oliva A. Resolving human object recognition in space and time. *Nat Neurosci* **17**, 455-
1170 462 (2014).

1171
1172 60. Marti S, Dehaene S. Discrete and continuous mechanisms of temporal selection in rapid visual streams.
1173 *Nature communications* **8**, 1955 (2017).

1174
1175 61. Quentin R*, et al.* Differential Brain Mechanisms of Selection and Maintenance of Information during Working
1176 Memory. *J Neurosci* **39**, 3728-3740 (2019).

1177
1178 62. Hebart MN, Bankson BB, Harel A, Baker CI, Cichy RM. The representational dynamics of task and object
1179 processing in humans. *Elife* **7**, (2018).

1180
1181 63. Grootswagers TR, A. K.; Shatek, S. M.; Carlson, T. A. . The neural dynamics underlying prioritisation of task-
1182 relevant information. *arXiv:210201303v2* (2021).

1183
1184 64. King JR, Dehaene S. Characterizing the dynamics of mental representations: the temporal generalization
1185 method. *Trends Cogn Sci* **18**, 203-210 (2014).

1186
1187 65. Chauhan T, Perales E, Xiao K, Hird E, Karatzas D, Wuerger S. The achromatic locus: effect of navigation
1188 direction in color space. *J Vis* **14**, (2014).

1189
1190 66. Pearce B, Crichton S, Mackiewicz M, Finlayson GD, Hurlbert A. Chromatic illumination discrimination ability
1191 reveals that human colour constancy is optimised for blue daylight illuminations. *PLoS One* **9**, e87989 (2014).

1192

1193 67. Lafer-Sousa R, Hermann KL, Conway BR. Striking individual differences in color perception uncovered by 'the
1194 dress' photograph. *Curr Biol* **25**, R545-546 (2015).

1195
1196 68. Winkler AD, Spillmann L, Werner JS, Webster MA. Asymmetries in blue-yellow color perception and in the
1197 color of 'the dress'. *Curr Biol* **25**, R547-548 (2015).

1198
1199 69. Gegenfurtner KR, Bloj M, Toscani M. The many colours of 'the dress'. *Curr Biol* **25**, R543-R544 (2015).

1200
1201 70. Brainard DH, Hurlbert AC. Colour Vision: Understanding #TheDress. *Curr Biol* **25**, R551-554 (2015).

1202
1203 71. Goddard E. A step toward understanding the human ventral visual pathway. *J Neurophysiol* **117**, 872-875
1204 (2017).

1205
1206 72. Lafer-Sousa R, Conway BR, Kanwisher NG. Color-Biased Regions of the Ventral Visual Pathway Lie between
1207 Face- and Place-Selective Regions in Humans, as in Macaques. *J Neurosci* **36**, 1682-1697 (2016).

1208
1209 73. Desikan RS*, et al.* An automated labeling system for subdividing the human cerebral cortex on MRI scans into
1210 gyral based regions of interest. *NeuroImage* **31**, 968-980 (2006).

1211
1212 74. Toga AW, Thompson PM, Mori S, Amunts K, Zilles K. Towards multimodal atlases of the human brain. *Nat Rev*
1213 *Neurosci* **7**, 952-966 (2006).

1214
1215 75. Duyck M*, et al.* Color Tuning of Face-Selective Neurons in Macaque Inferior Temporal Cortex. *Eneuro* **8**,
1216 (2021).

1217
1218 76. Hermann K, Pantazis D, Conway BR. The dynamics of color processing in humans measured with MEG. *Society*
1219 *for Neuroscience Annual Meeting 79003/M32*, (2015).

1220
1221 77. Rosenthal IA, Hermann KL, Vonder Haar C, Pantazis D, Conway BR. Decoding hue and luminance with
1222 magnetoencephalography Society for Neuroscience annual meeting, 774.03. (2017).

1223
1224 78. Teichmann L, Grootswagers T, Carlson TA, Rich AN. Seeing versus knowing: The temporal dynamics of real
1225 and implied colour processing in the human brain. *NeuroImage* **200**, 373-381 (2019).

1226
1227 79. Sandhaeger F, von Nicolai C, Miller EK, Siegel M. Monkey EEG links neuronal color and motion information
1228 across species and scales. *Elife* **8**, (2019).

1229
1230 80. Hajonides JE, Nobre AC, van Ede F, Stokes MG. Decoding visual colour from scalp electroencephalography
1231 measurements. *BioRxiv* **https://doi.org/10.1101/2020.07.30.228437** (2020).

1232
1233 81. Chauhan T, Jakovljev, I., Thompson, L. N., Wuerger, S. M., Martinovic, J. Decoding of EEG signals reveals non-
1234 uniformities in the neural geometry of colour. *bioRxiv* **2021.06.17.448044**, (2021).

82. Xiao Y, Wang Y, Felleman DJ. A spatially organized representation of colour in macaque cortical area V2. *Nature* **421**, 535-539 (2003).

83. Xiao Y, Casti A, Xiao J, Kaplan E. Hue maps in primate striate cortex. *NeuroImage* **35**, 771-786 (2007).

84. Garg AK, Li P, Rashid MS, Callaway EM. Color and orientation are jointly coded and spatially organized in primate primary visual cortex. *Science* **364**, 1275-1279 (2019).

85. Liu Y*, et al.* Hierarchical Representation for Chromatic Processing across Macaque V1, V2, and V4. *Neuron*, (2020).

86. Conway BR, Tsao DY. Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex. *Proceedings of the National Academy of Science (USA)* **106**, 18034-18039 (2009).

87. Stokes MG, Wolff MJ, Spaak E. Decoding Rich Spatial Information with High Temporal Resolution. *Trends Cogn Sci* **19**, 636-638 (2015).

88. Hermann K, Rosenthal I, Singh S, Pantazis D, Conway BR. Temporal dynamics of the neural mechanisms for encoding hue and luminance contrast uncovered by magnetoencephalography. *BioRxiv* https://www.biorxiv.org/content/10.1101/2020.06.17.155713v2, (2020).

89. Sutterer DW, Coia AJ, Sun V, Shevell SK, Awh E. Decoding chromaticity and luminance from patterns of EEG activity. *Psychophysiology* **58**, e13779 (2021).

90. Berninger TA, Arden GB, Hogg CR, Frumkes T. Separable evoked retinal and cortical potentials from each major visual pathway: preliminary results. *Br J Ophthalmol* **73**, 502-511 (1989).

91. Kulikowski JJ, Robson AG, Murray IJ. Scalp VEPs and intra-cortical responses to chromatic and achromatic stimuli in primates. *Doc Ophthalmol* **105**, 243-279 (2002).

92. Zhang W, Luck SJ. Feature-based attention modulates feedforward visual processing. *Nat Neurosci* **12**, 24-25 (2009).

93. Proklova D, Kaiser D, Peelen MV. MEG sensor patterns reflect perceptual but not categorical similarity of animate and inanimate objects. *NeuroImage* **193**, 167-177 (2019).

94. Carlson T, Tovar DA, Alink A, Kriegeskorte N. Representational dynamics of object vision: the first 1000 ms. *J Vis* **13**, (2013).

95. Kaiser D, Azzalini DC, Peelen MV. Shape-independent object category responses revealed by MEG and fMRI decoding. *J Neurophysiol* **115**, 2246-2250 (2016).

96. Dobs K, Isik L, Pantazis D, Kanwisher N. How face perception unfolds over time. *Nature communications* **10**, 1258 (2019).

97. Lafer-Sousa R, Conway BR. Parallel, multi-stage processing of colors, faces and shapes in macaque inferior temporal cortex. *Nat Neurosci* **16**, 1870-1878 (2013).

98. Bannert MM, Bartels A. Human V4 Activity Patterns Predict Behavioral Performance in Imagery of Object Color. *Journal of Neuroscience* **38**, 3657-3668 (2018).

99. Siuda-Krzywicka K, Witzel C, Bartolomeo P, Cohen L. Color Naming and Categorization Depend on Distinct Functional Brain Networks. *Cereb Cortex*, (2020).

100. Conway BR. Color vision, cones, and color-coding in the cortex. *Neuroscientist* **15**, 274-290 (2009).

101. Namima T, Yasuda M, Banno T, Okazawa G, Komatsu H. 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).

102. Crognale MA, Duncan CS, Shoenhard H, Peterson DJ, Berryhill ME. The locus of color sensation: cortical color loss and the chromatic visual evoked potential. *J Vis* **13**, (2013).

103. Cauchoix M, Barragan-Jason G, Serre T, Barbeau EJ. The neural dynamics of face detection in the wild revealed by MVPA. *J Neurosci* **34**, 846-854 (2014).

104. Goddard E, Carlson TA, Dermody N, Woolgar A. Representational dynamics of object recognition: Feedforward and feedback information flows. *NeuroImage* **128**, 385-397 (2016).

105. Martin Cichy R, Khosla A, Pantazis D, Oliva A. Dynamics of scene representations in the human brain revealed by magnetoencephalography and deep neural networks. *NeuroImage* **153**, 346-358 (2017).

106. Rathbun DL, Warland DK, Usrey WM. Spike timing and information transmission at retinogeniculate synapses. *J Neurosci* **30**, 13558-13566 (2010).

107. King J-R, Wyart V. The Human Brain encodes a Chronicle of Visual Events at each Instant of Time. *bioRxiv*, 846576 (2019).

108. Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology* **195**, 215-243 (1968).

109. Horwitz GD. Signals Related to Color in the Early Visual Cortex. *Annu Rev Vis Sci* **6**, 287-311 (2020).

110. Lafer-Sousa R, Conway BR. #TheDress: Categorical perception of an ambiguous color image. *J Vis* **17**, 25 (2017).

111. Cicmil N, Bridge H, Parker AJ, Woolrich MW, Krug K. Localization of MEG human brain responses to retinotopic visual stimuli with contrasting source reconstruction approaches. *Front Neurosci* **8**, 127 (2014).

112. Westland S, Ripamonti C, Cheung V. Computational colour science using matlab (2nd Edition). *John Wiley and Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, England*, (2012).

113. Brainard DH. Cone contrast and opponent modulation color spaces in *Human Color Vision* (P.K. Kaiser and R.M. Boynton Editors, 2nd Edition). *Washington, DC: Optical Society of America*, (1996).

114. Kleiner M, Brainard DH, Pelli D. What's new in Psychtoolbox-3? *Perception* **36**, (2007).

115. Isik L, Meyers EM, Leibo JZ, Poggio T. The dynamics of invariant object recognition in the human visual system. *J Neurophysiol* **111**, 91-102 (2014).

116. Maris E, Oostenveld R. Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* **164**, 177-190 (2007).

Figure Legends

Figure 1. Decoding luminance polarity and hue from MEG data: approach

a) Decoding luminance polarity. Participants were shown four hues that varied in luminance polarity (light/dark). Classifiers were trained (solid arrows) and tested (dashed arrows) using stimuli of the same hue (identity problems) or different hue (generalization problems). to determine the extent to which the MEG response to a given color is informative of the luminance polarity carried by the same hue (4 identity problems) or by other hues (12 generalizing-across-hue problems). Each binary classifier was trained to distinguish whether a light or dark stimulus had been presented given patterns of magnetoencephalography (MEG) sensor activations. For the four identity problems, classifiers were trained and tested on the same hue (all four identity problems are illustrated in the graphic). For the generalization problems, classifiers were trained and tested on different hues (half of the 12 generalizing-across-hue problems are illustrated). In the graphic, solid lines indicate comparisons used for training and dashed lines indicate comparisons used for testing; line shading distinguishes different problems. **b)** Decoding

hue. Format as in panel (a). For stimuli of a given luminance polarity (e.g. dark), a binary classifier was trained to determine which of two hues (e.g. pink or orange) had been presented. The classifiers were then tested on held-out trials in which the luminance polarity (e.g. dark) was the same as at train time (12 identity problems) or in which the luminance polarity was opposite (e.g. light), requiring generalization of hue across luminance polarity (12 generalizing-across-luminance-polarity problems).

Figure 2. Experimental paradigm

Participants were scanned with magnetoencephalography (MEG) while they looked at colored spirals that were flashed on the screen for 116 ms (with 1 second between flashes). The spirals were one of eight colors: four hues of either a luminance increment or luminance decrement (**Supplementary Figure 1** provides the color specifications for the stimuli). Stimuli were pseudo-randomly interleaved, with 500 presentations of each stimulus over two recording sessions for each of 18 participants. Trials during eye blinks or other artifacts were removed, and the remaining trials were randomly subsampled to yield 375 trials per condition. Sensor data were averaged into 5ms bins, 200ms before stimulus onset to 600ms after stimulus onset. The figure shows simulated data to illustrate the analysis pipeline. Classifiers were trained at every time point, and individually for each participant: at each time point (*t*) in the 800ms time window, the 375 trials were divided into 5 sets of 75. Four sets were used to train the classifier, and 1 set was used to test the classifier. The procedure was repeated for the 5 cross-validation splits; and the entire procedure was repeated 50x with different random assignments of the 375 trials into the 5 sets, yielding a decoding curve showing classification accuracy as a function of time after stimulus onset (bottom right showing actual data for classifying luminance polarity using pink). The horizontal sequence of data points above the x-axis shows the time points when decoding was significant for more than 5 consecutive time points (FDR corrected; the gray arrow shows the classification accuracy and the time point at which decoding became significant).

Figure 3. Decoding luminance polarity and hue from MEG data: results

a) Classification accuracy as a function of time after stimulus onset for the luminance-polarity problems (average of identity problems, solid line; average of generalization problems, dashed line). Traces were generated by averaging 1000 bootstrapped samples across 18 participants. Shading shows the SE of the bootstrap samples. The stimulus duration was 116 ms (gray bar). Inset shows the difference in peak for the 1000 bootstrapped comparisons: identity minus generalization (mean = -6 ms, p=0.34; one-sided test of the proportion of these differences that was less than or equal to zero, not corrected for multiple comparisons). Initial peak of the identity problems, solid vertical gray line, 104.4 ms [100 110]; initial peak of the generalization problems, dashed gray line, 96.5 ms [85 100]. Open arrowhead shows the second decoding peak, which corresponds to stimulus cessation. The horizontal sequence of data points above the x-axis, demarcated by asterisks, show time points at which decoding was above chance (determined by a permutation test across subjects and cluster corrected, see Methods); onset of significant decoding for the identity problems was achieved at 65 ms [60, 80], for the generalization problems, at 55 ms [50, 70]. **b)** The average performance across 12 sets of identity problems (solid line) and 12 sets of generalization problems (dashed line). The inset shows the difference in peak (for the 1000 bootstrapped comparisons across participants, identity minus generalization, mean = 1 ms, p=0.58, proportion of differences that was less than or equal to zero, not corrected for multiple comparisons). The time to peak was the same for the identity and generalization problems (identity problems: solid vertical line, 121.8 ms [120, 130]; generalization problems: dashed vertical line, 119.1 ms [115, 130]). Onset of significant decoding for the identity problems was achieved at 80 ms [75, 85], and for the generalization problems, at 115 [35, 355]. Other conventions as for panel (b).

Figure 4. Comparing the temporal dynamics of decoding, luminance polarity versus hue

a) Classification accuracy for the identity problems (see Figure 3a,c). The data used to decode luminance identity was subsampled so that the maximum classification accuracy for luminance identity was the same as that of hue identity (61% of the trials were used). Luminance polarity decoding peaked at 103.4 ms [100, 110]. Inset shows the differences between the peaks across the 1000 bootstrapped samples across the 18 participants. Accuracy for

luminance polarity peaked 18 ms before hue (p=0.003). The magnitude of classifier accuracy corresponding to the cessation of the stimulus (open arrowhead) was greater for luminance polarity than hue (p<0.001; exact p values not possible, see Methods; not adjusted for multiple comparisons). **b)** Classification accuracy for the generalization problems. The data used to decode the generalizing luminance problem was subsampled at 16% of the original data to match the initial peak of generalizing hue. The subsampled luminance polarity problem peaked at 99.8 ms [95,110]. Accuracy for luminance polarity peaked 19 ms before hue (p=0.017). The magnitude of classifier accuracy corresponding to the cessation of the stimulus was not greater for luminance polarity than hue (p=0.3). The p values in a and b come from a one-sided statistical test comparing the bootstrapped distribution of the times to peak. The p value is the proportion of the 1000 bootstrapped times to peak of the hue problems that were earlier or equal to the bootstrapped times to peak of the luminance problems. The p values are not corrected for multiple comparisons. Other conventions as Figure 3 b,d.

Figure 5. Task had no impact on the time to peak decoding of hue or luminance polarity. a) Tasks, 1-back luminance-polarity matching (left), 1-back hue matching task (right). **b) Left panel,** Decoding luminance polarity (black trace) and hue (blue trace) using data from the luminance-polarity matching task. The data show the average of the identity decoding problems (see Figure 1). Traces represent the average of 1000 bootrapped time courses across 8 sessions (four sessions completed by two subjects). Shading is SEM of the bootstrapped time courses. Vertical lines show the time to peak. The horizontal sequence of points above the x-axis show when decoding was greater than chance for five or more consecutive time bins (determined using within-subjects null decoding distribution and FDR corrected, see Methods). Hue decoding peaked after luminance-polarity decoding (p<0.001). **Right panel,** Decoding luminance polarity (black line) and hue (blue line) using data from the hue-matching task. Hue decoding peaked after luminance-polarity decoding (p<0.001). Other conventions as for the left panel. **c), Left panel,** Decoding luminance polarity using data from the luminance-polarity matching task (bold line) and the hue-matching task (thin line). Shading represents the SEM. The time-to-peak for decoding luminance polarity identity was not impacted by task (p=0.384). **Right panel,** Decoding hue using data from the hue-matching task (bold line) and luminance-polarity matching task (thin line); the time-to-peak for decoding hue was not impacted by task (p=0.427; p values reflect two-sided t tests using a bootstrapping procedure, not adjusted for multiple comparisons). The bolded traces show the decoding problems that aligned with the task. Luminance-polarity decoding was significant for a longer duration for data collected during the luminance-matching task (left panel), while hue decoding was significant for a longer duration for data collected during the hue-matching task (right panel).

Figure 6. Cross-temporal generalization of decoding, luminance polarity versus hue

a) Classifiers were trained using the pattern of magnetoencephalography activity elicited at time points from -200 ms to 600 ms after stimulus onset (y-axis) and tested using data not used in the training, across the same time interval, on the set of problems for decoding luminance polarity generalizing across hue. The best decoding performance was achieved for classifiers that were trained and tested using data from the same time point after stimulus onset, indicated by the strong performance along the x=y diagonal. The peak classification time was at 100 ms; there was a dip in classification performance at 130 ms. The black contours show regions in the heatmap that were p<0.05 cluster corrected. The p values were obtained using a sign-permutation test. **b**) Data as in panel (a), but for classifiers trained and tested on the set of problems for decoding hue generalizing across luminance polarity. The peak classification time was 119 ms; the cluster-corrected regions of significance extended further from the diagonal compared to panel (a). **c)** Comparison of results in panel (a) and (b). The time points where the classifiers were more accurate for luminance polarity compared to hue are shown as dark blue (white contours show cluster-corrected significant results), while the time points where the classifiers were more accurate for hue compared to luminance polarity are shown as yellow (red contours show cluster-corrected results). The p values were obtained using a permutation test

on the difference in accuracy between hue and luminance decoding. The p values are cluster corrected. Decoding of hue showed greater generalization across time compared to decoding of luminance polarity, and the greater cross-temporal generalization for hue began relatively early after stimulus onset (~119 ms).

Figure 7. Decoding luminance polarity carried by different colors

a) Accuracy of classifiers trained to decode luminance polarity on the identity problems (individual problems, thin lines; average, thick line). **b)** Accuracy of classifiers trained to decode luminance polarity across hues (12 generalizing-across-hue problems, thin lines; average, thick dashed line; shading shows SE across subjects). The thickness of the horizontal line above the x-axis shows the number of decoding problems ((a) of 4; (b) of 12) that were significant in each 5ms bin (cluster corrected). **c)** Heatmap showing classifier performance [95% CI, by bootstrapping] for each binary classifier at the initial peak decoding for each problem**.** Bold numbers show the total number of 5ms time bins with significant classification performance. The time stamps show when the classifier performance reached the initial peak in decoding accuracy. **d)** Duration of significant decoding of luminance polarity plotted in the cone-opponent color space used to define the stimuli; B, P, O, and G demarcate within this space the location of the blue, pink, orange, and green hues of the stimuli. The color scalebar indicates the number of 5ms time bins that were significant, and the arrow direction indicates the hues of the stimuli for training and testing. For example, the arrow from P to G corresponds to classifiers trained on luminance polarity using data obtained with light/dark pink stimuli and tested using data obtained with light/dark green stimuli. For the four identity problems (the letters B, P, O, and G), the hue was the same for training and testing. For each type of problem, the precise number of significant time bins is shown, along with the [95% CI] obtained by bootstrapping. Problems that cross the color space, in which the hues for training and testing varied in both L-M and S activation, showed lower numbers of significant time bins than problems around the perimeter of the color space, in which the hues for training and testing varied in only L-M or S.

Figure 8. Source localization to regions defined using fMRI

a) Functional anatomy from one participant, on the inflated cortical surface (left panel, side view, anterior to left; right panel, ventral view, anterior at top; similar results were obtained in 14 participants). Regions of interest were defined using fMRI responses to movie clips, in color and black-and-white, of faces, bodies, objects, scrambled objects.[68] The activation map (yellow-orange) shows voxels with higher responses to color clips compared to black-and-white clips. Contours show ROIs for faces > objects (including the fusiform face area, FFA), intact shapes > scrambled shapes (lateral object area, LO), and places > objects (including parahippocampal place area, PPA). Regions of interest for V1, V2, MT (middle temporal area), frontal, and precentral ROIs were defined using an anatomical atlas. **b)** Source localization estimates of the current source density (CSD) of magnetoencephalography (MEG) responses to color arising from each functional ROI, averaged across participants and calculated with dynamical Statistical Parametric Mapping (dSPM; see Methods). 0 on the x-axis is stimulus onset. Values on the y-axis are unitless. Transparent shading shows SEM. **c)** Maximum amplitude of CSD in each ROI of the ventral visual pathway, calculated as distance from peak to trough of the time course in panel b. Error bars are SEM (n=14). There was a significant effect of ROI on response magnitude (repeated measures one-way ANOVA, p = 0.002). Responses source-localized to the color-biased regions were different from those of LO (two-sided paired t-test, p=0.01) and place-biased regions (p=0.02), but not face-biased ROIs (p=0.12). **d)** Average classifier performance on decoding luminance contrast generalizing-across-hue (12 problems averaged together; see Figure 1a) trained using only those MEG data localized to the MRI-defined ROIs (N=14 participants). Each line shows the average accuracy of one ROI-restricted classifier averaged across participants (color key in panel a). **e)** Average classifier performance on decoding hue generalizing-across-luminance contrast (12 problems averaged together; see Figure 1b). **f)** The distribution of sensors used as features for the classifiers across participants (N=18). Color bar shows the percent likelihood that any given sensor was selected as a feature. **g)** As in panel (f), but for decoding hue generalizing across luminance contrast.

1498

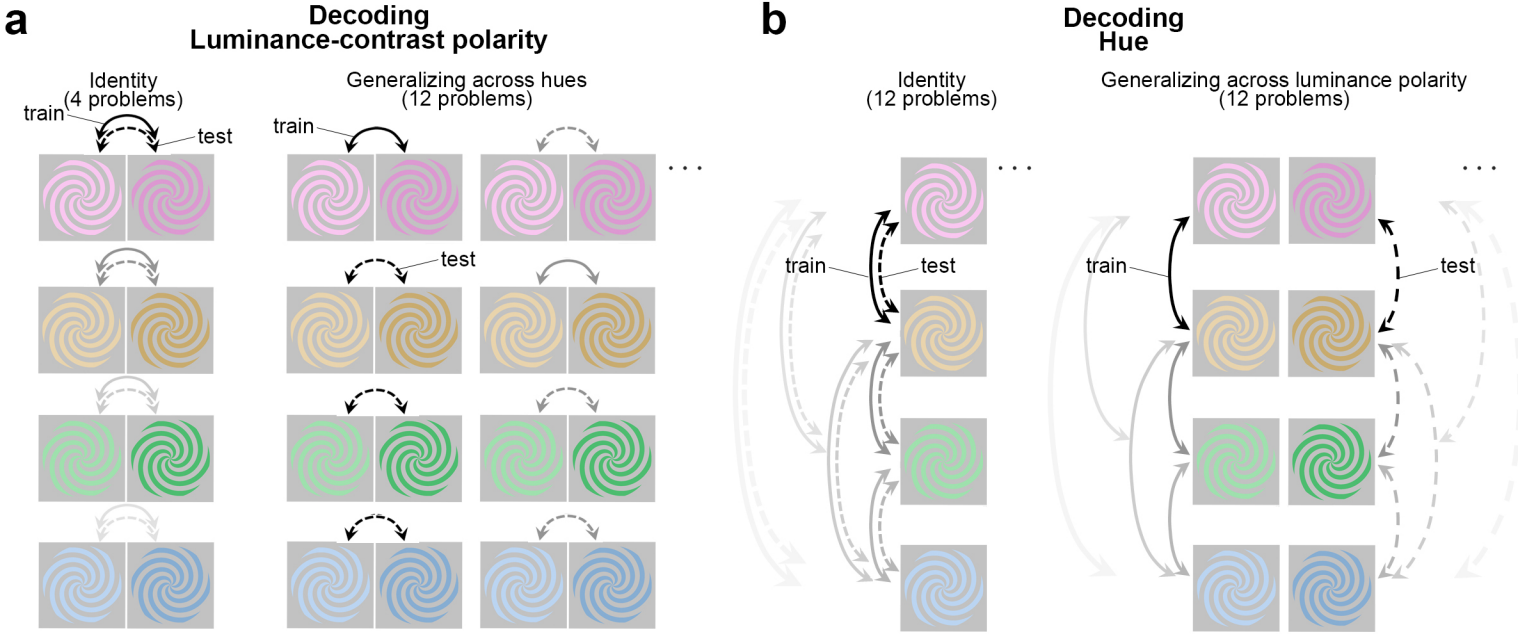
a
Decoding
Luminance-contrast polarity
Identity
(4 problems)
train
test
Generalizing across hues
(12 problems)
train
test
b
Decoding
Hue
Identity
(12 problems)
train
test
Generalizing across luminance polarity
(12 problems)
train
test

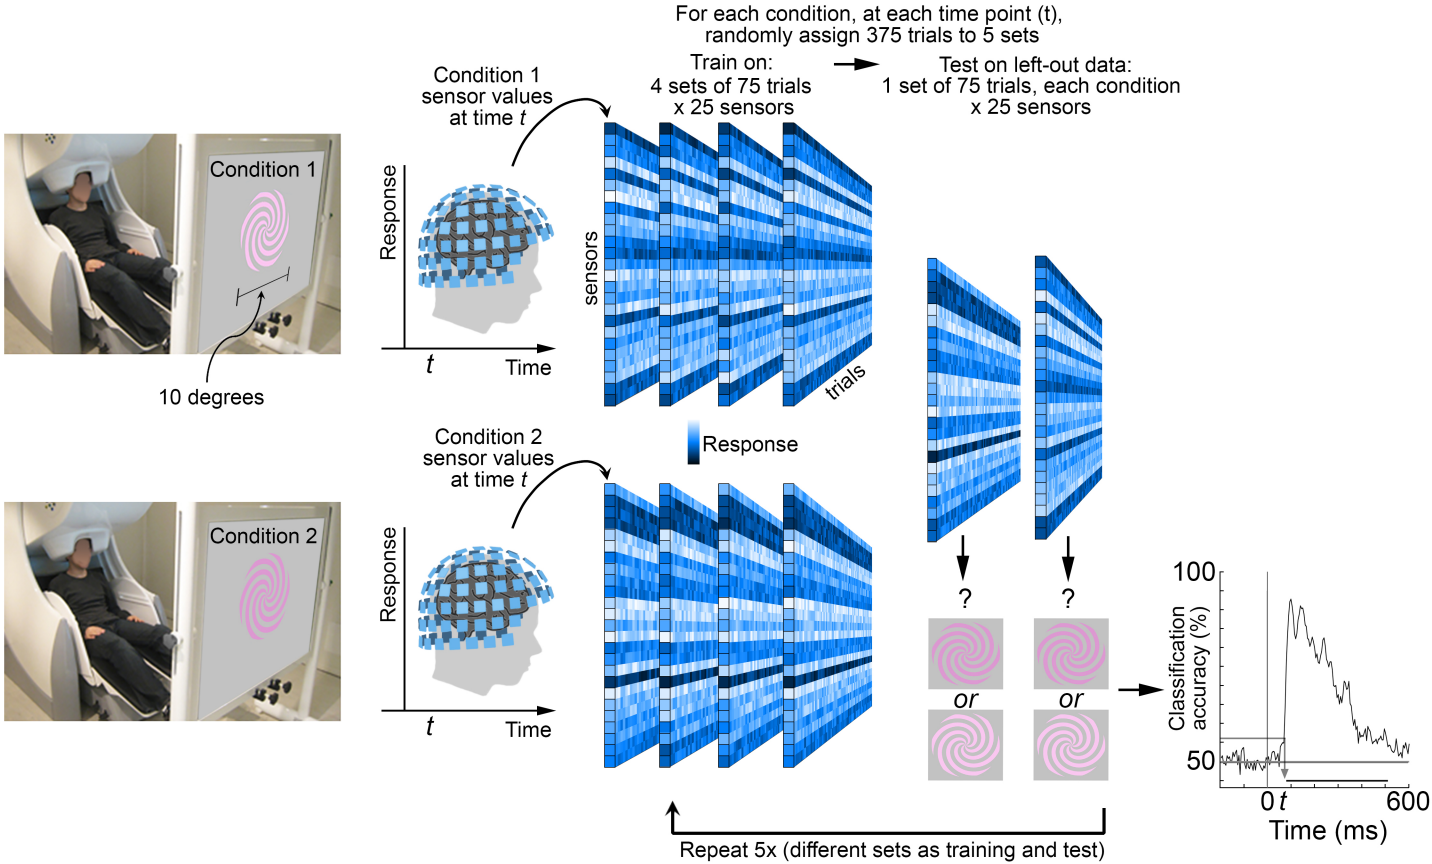
For each condition, at each time point (t),
randomly assign 375 trials to 5 sets
Train on:
4 sets of 75 trials
x 25 sensors
Test on left-out data:
1 set of 75 trials, each condition
x 25 sensors
Condition 1
Condition 1
sensor values
at time t
Response
t
Time
10 degrees
sensors
trials
Response
Condition 2
Condition 2
sensor values
at time t
Response
t
Time
?
?
or
or
100
Classification
accuracy (%)
50
0 t
600
Time (ms)
Repeat 5x (different sets as training and test)

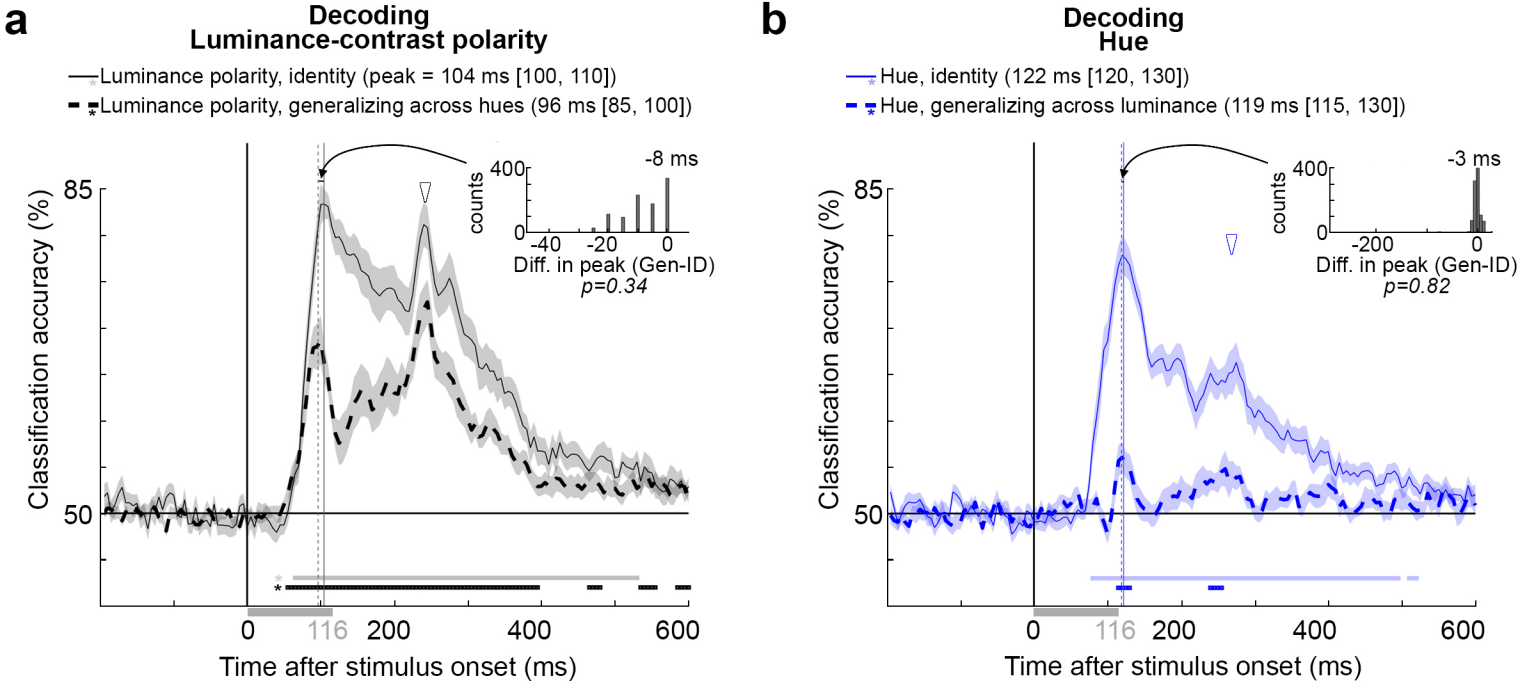
a
Decoding
Luminance-contrast polarity
Luminance polarity, identity (peak = 104 ms [100, 110])
Luminance polarity, generalizing across hues (96 ms [85, 100])
Classification accuracy (%)
85
50
0
116
200
400
600
Time after stimulus onset (ms)
counts
400
0
-8 ms
-40
-20
0
Diff. in peak (Gen-ID)
p=0.34
b
Decoding
Hue
Hue, identity (122 ms [120, 130])
Hue, generalizing across luminance (119 ms [115, 130])
Classification accuracy (%)
85
50
0
116
200
400
600
Time after stimulus onset (ms)
counts
400
0
-3 ms
-200
0
Diff. in peak (Gen-ID)
p=0.82

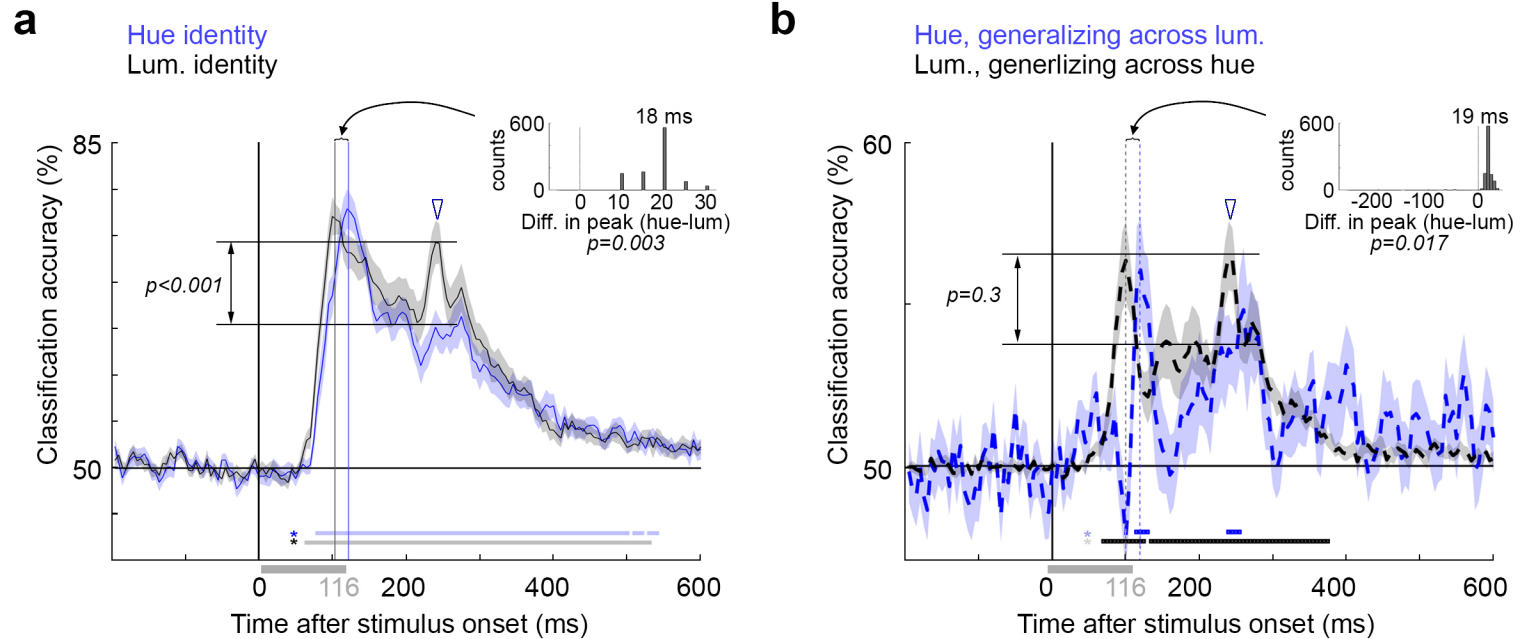
a
Hue identity
Lum. identity
Classification accuracy (%)
85
50
p<0.001
18 ms
counts
600
0
0 10 20 30
Diff. in peak (hue-lum)
p=0.003
0 116 200 400 600
Time after stimulus onset (ms)
b
Hue, generalizing across lum.
Lum., generlizing across hue
Classification accuracy (%)
60
50
p=0.3
19 ms
counts
600
0
-200 -100 0
Diff. in peak (hue-lum)
p=0.017
0 116 200 400 600
Time after stimulus onset (ms)

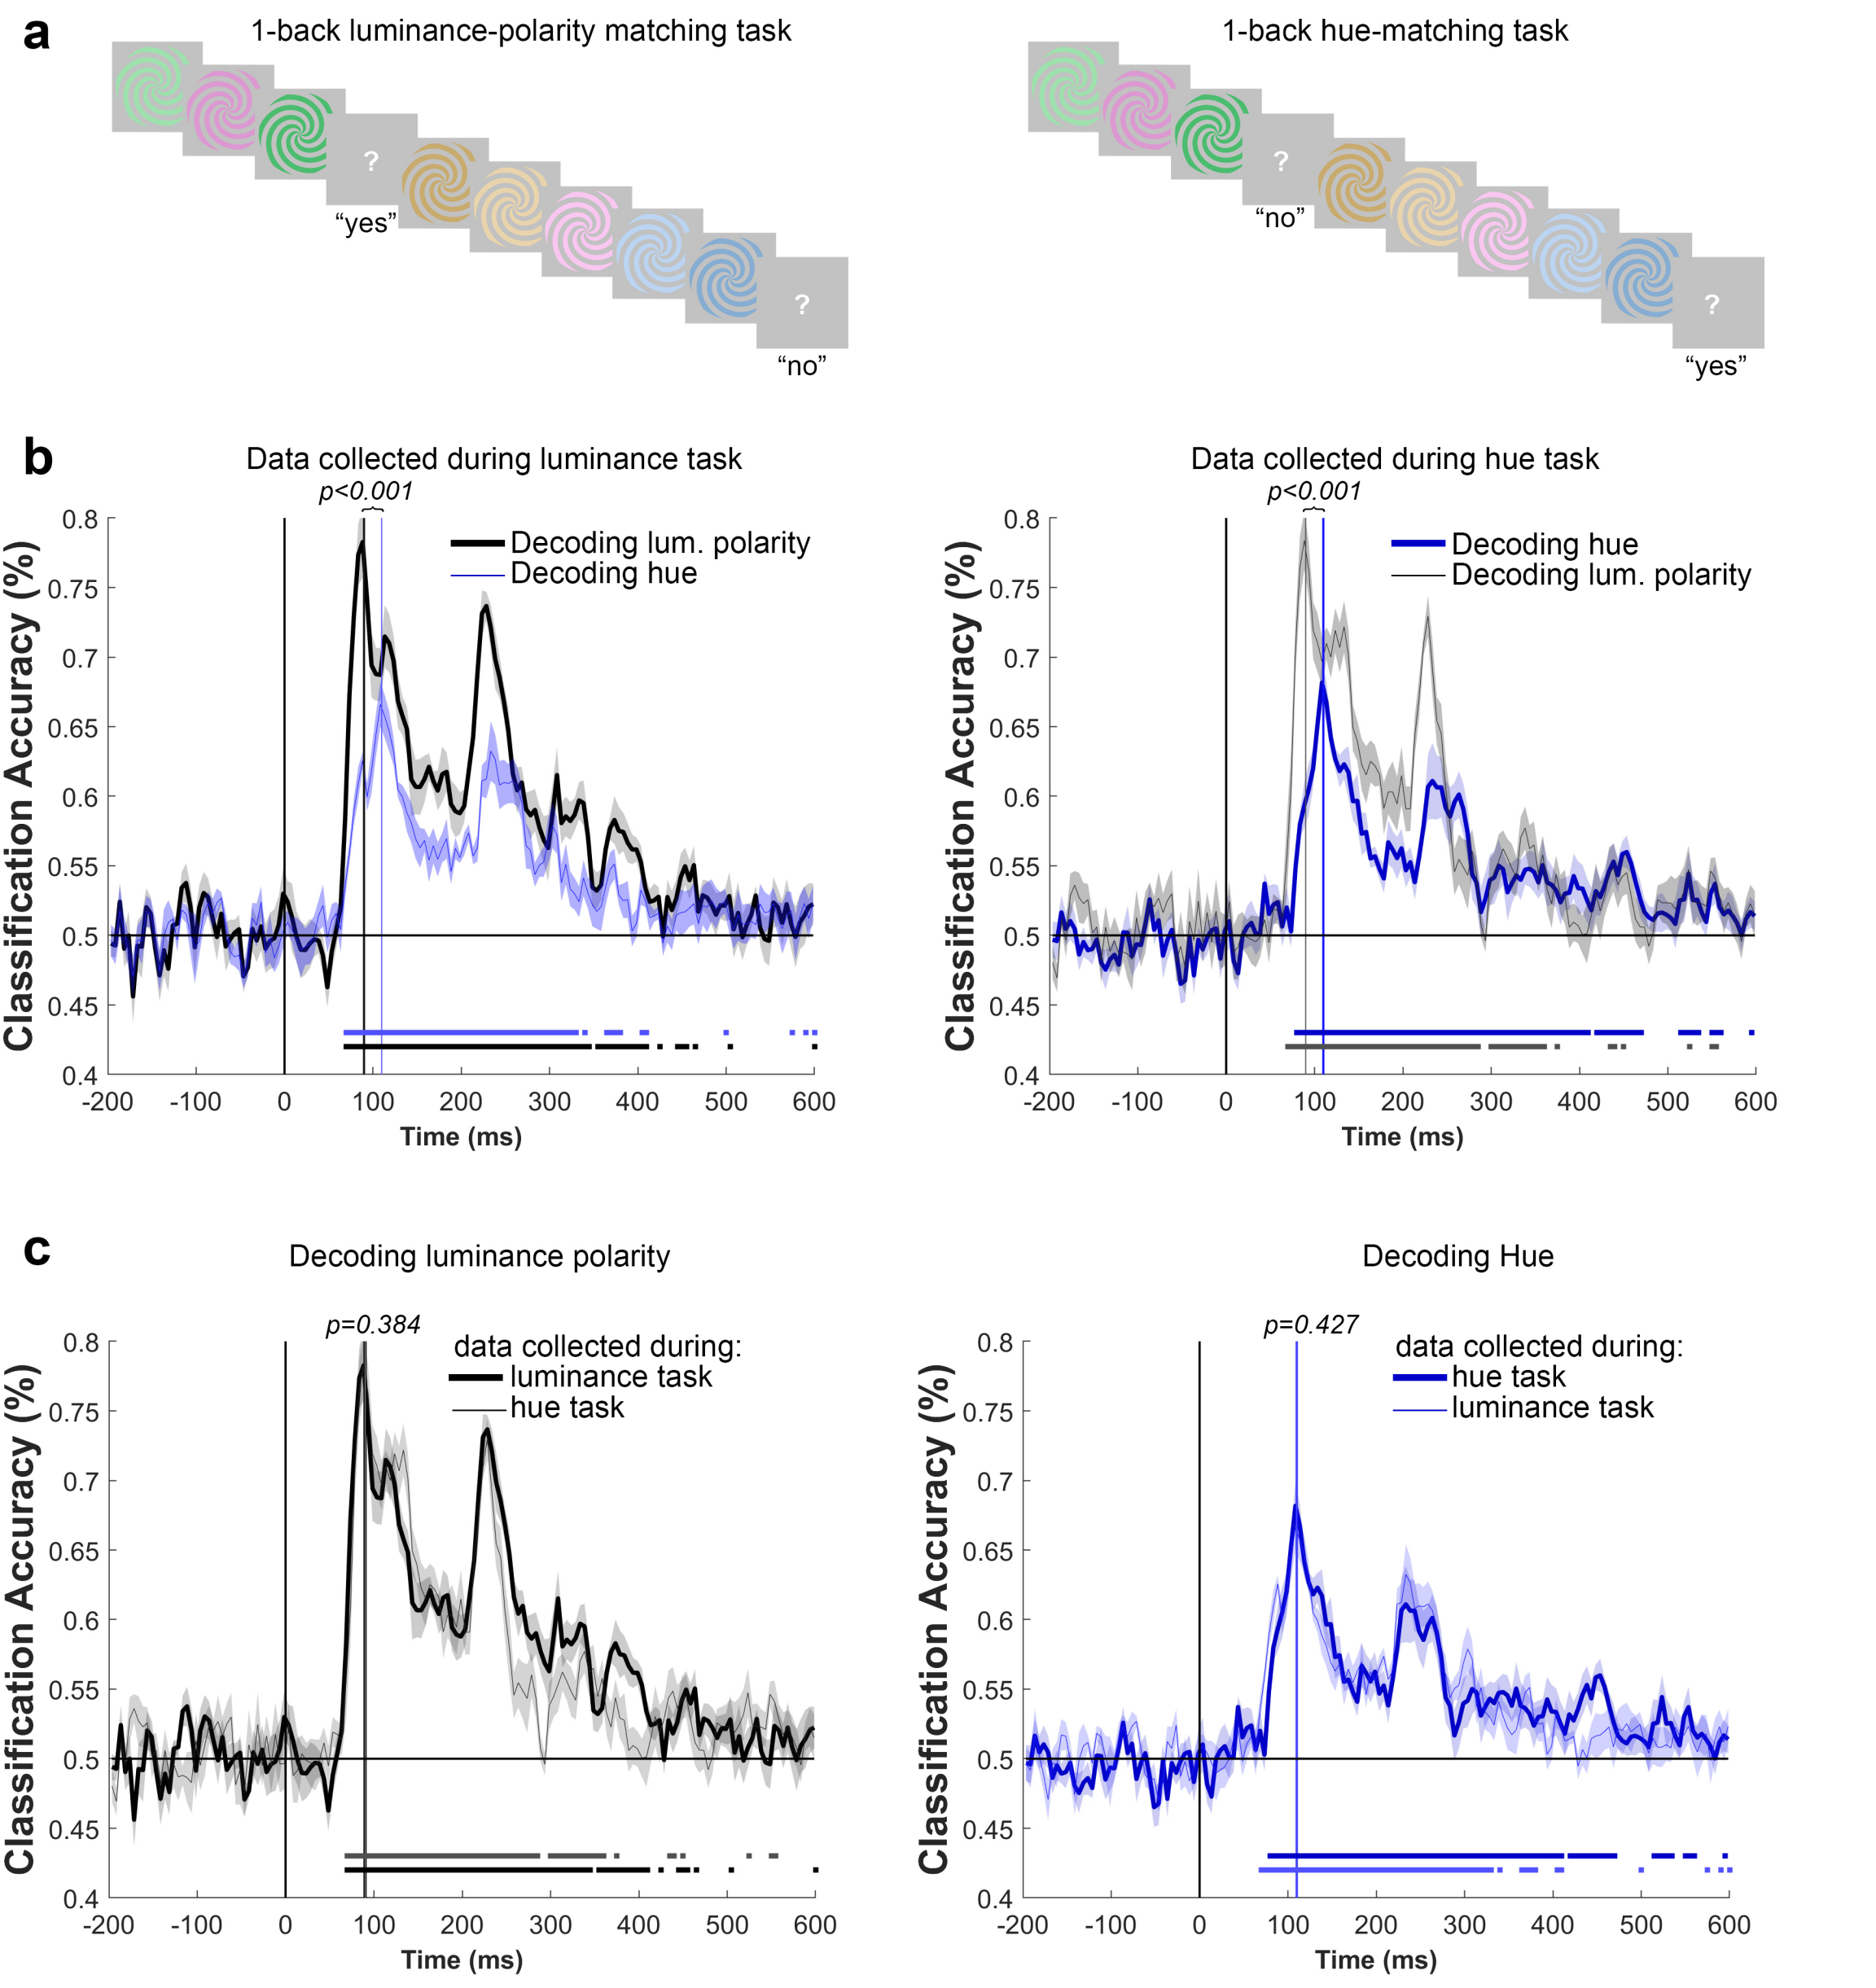
a
1-back luminance-polarity matching task
?
"yes"
?
"no"
1-back hue-matching task
?
"no"
?
"yes"
b
Data collected during luminance task
p<0.001
Decoding lum. polarity
Decoding hue
Classification Accuracy (%)
Time (ms)
Data collected during hue task
p<0.001
Decoding hue
Decoding lum. polarity
Classification Accuracy (%)
Time (ms)
c
Decoding luminance polarity
p=0.384
data collected during:
luminance task
hue task
Classification Accuracy (%)
Time (ms)
Decoding Hue
p=0.427
data collected during:
hue task
luminance task
Classification Accuracy (%)
Time (ms)

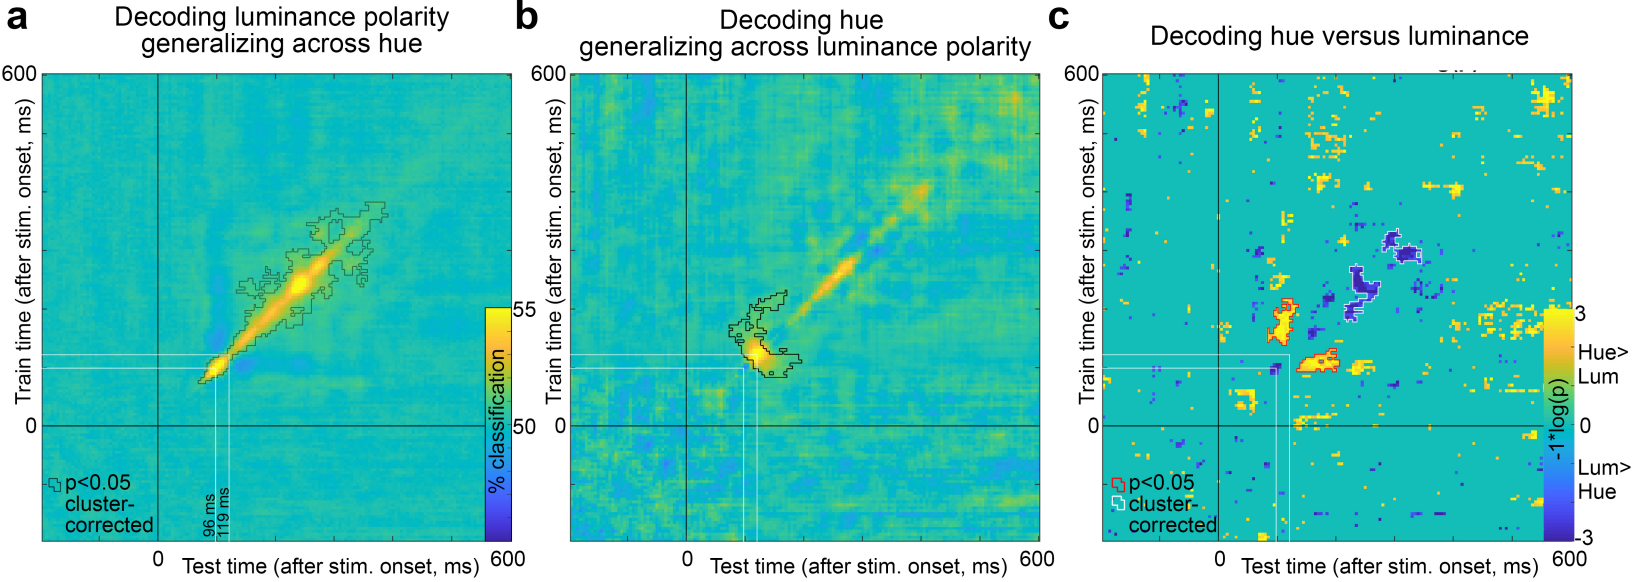
a
Decoding luminance polarity
generalizing across hue
b
Decoding hue
generalizing across luminance polarity
c
Decoding hue versus luminance
Train time (after stim. onset, ms)
Test time (after stim. onset, ms)
600
0
55
50
% classification
p<0.05
cluster-
corrected
96 ms
119 ms
3
-3
-1*log(p)
Hue>
Lum
Lum>
Hue

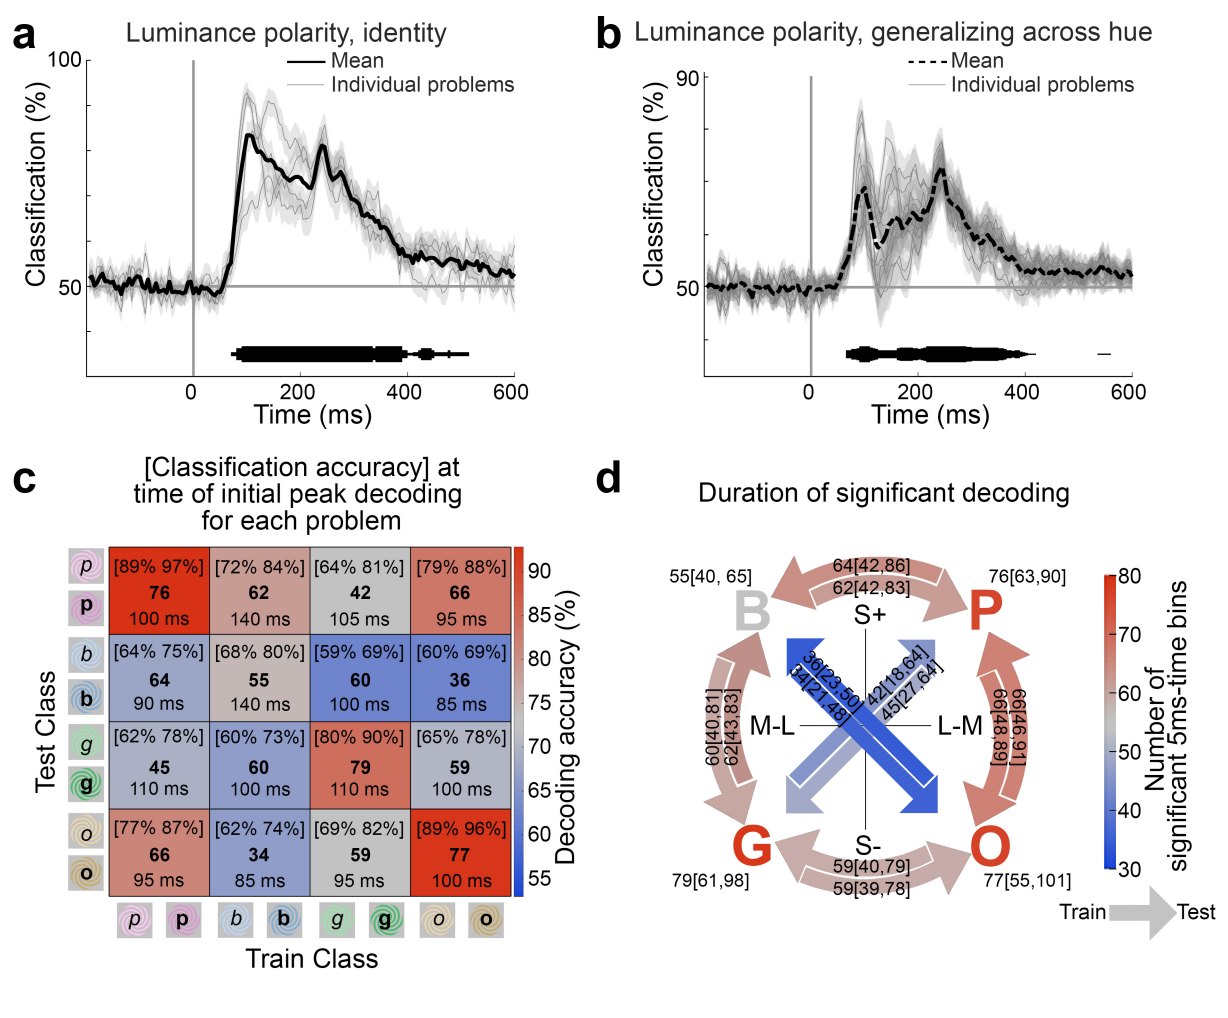
a
Luminance polarity, identity
Mean
Individual problems
Classification (%)
100
50
0
200
400
600
Time (ms)
b
Luminance polarity, generalizing across hue
Mean
Individual problems
Classification (%)
90
50
0
200
400
600
Time (ms)
c
[Classification accuracy] at time of initial peak decoding for each problem
Test Class
Train Class
Decoding accuracy (%)
90
85
80
75
70
65
60
55
p [89% 97%] 76 100 ms
p [72% 84%] 62 140 ms
p [64% 81%] 42 105 ms
p [79% 88%] 66 95 ms
b [64% 75%] 64 90 ms
b [68% 80%] 55 140 ms
b [59% 69%] 60 100 ms
b [60% 69%] 36 85 ms
g [62% 78%] 45 110 ms
g [60% 73%] 60 100 ms
g [80% 90%] 79 110 ms
g [65% 78%] 59 100 ms
o [77% 87%] 66 95 ms
o [62% 74%] 34 85 ms
o [69% 82%] 59 95 ms
o [89% 96%] 77 100 ms
d
Duration of significant decoding
B
P
G
O
S+
S-
M-L
L-M
55[40, 65]
64[42,86]
62[42,83]
76[63,90]
60[40,81]
62[43,83]
36[23,50]
34[21,48]
42[18,64]
45[27,64]
66[46,91]
66[48,88]
79[61,98]
59[40,79]
59[39,78]
77[55,101]
Number of significant 5ms-time bins
80
70
60
50
40
30
Train
Test

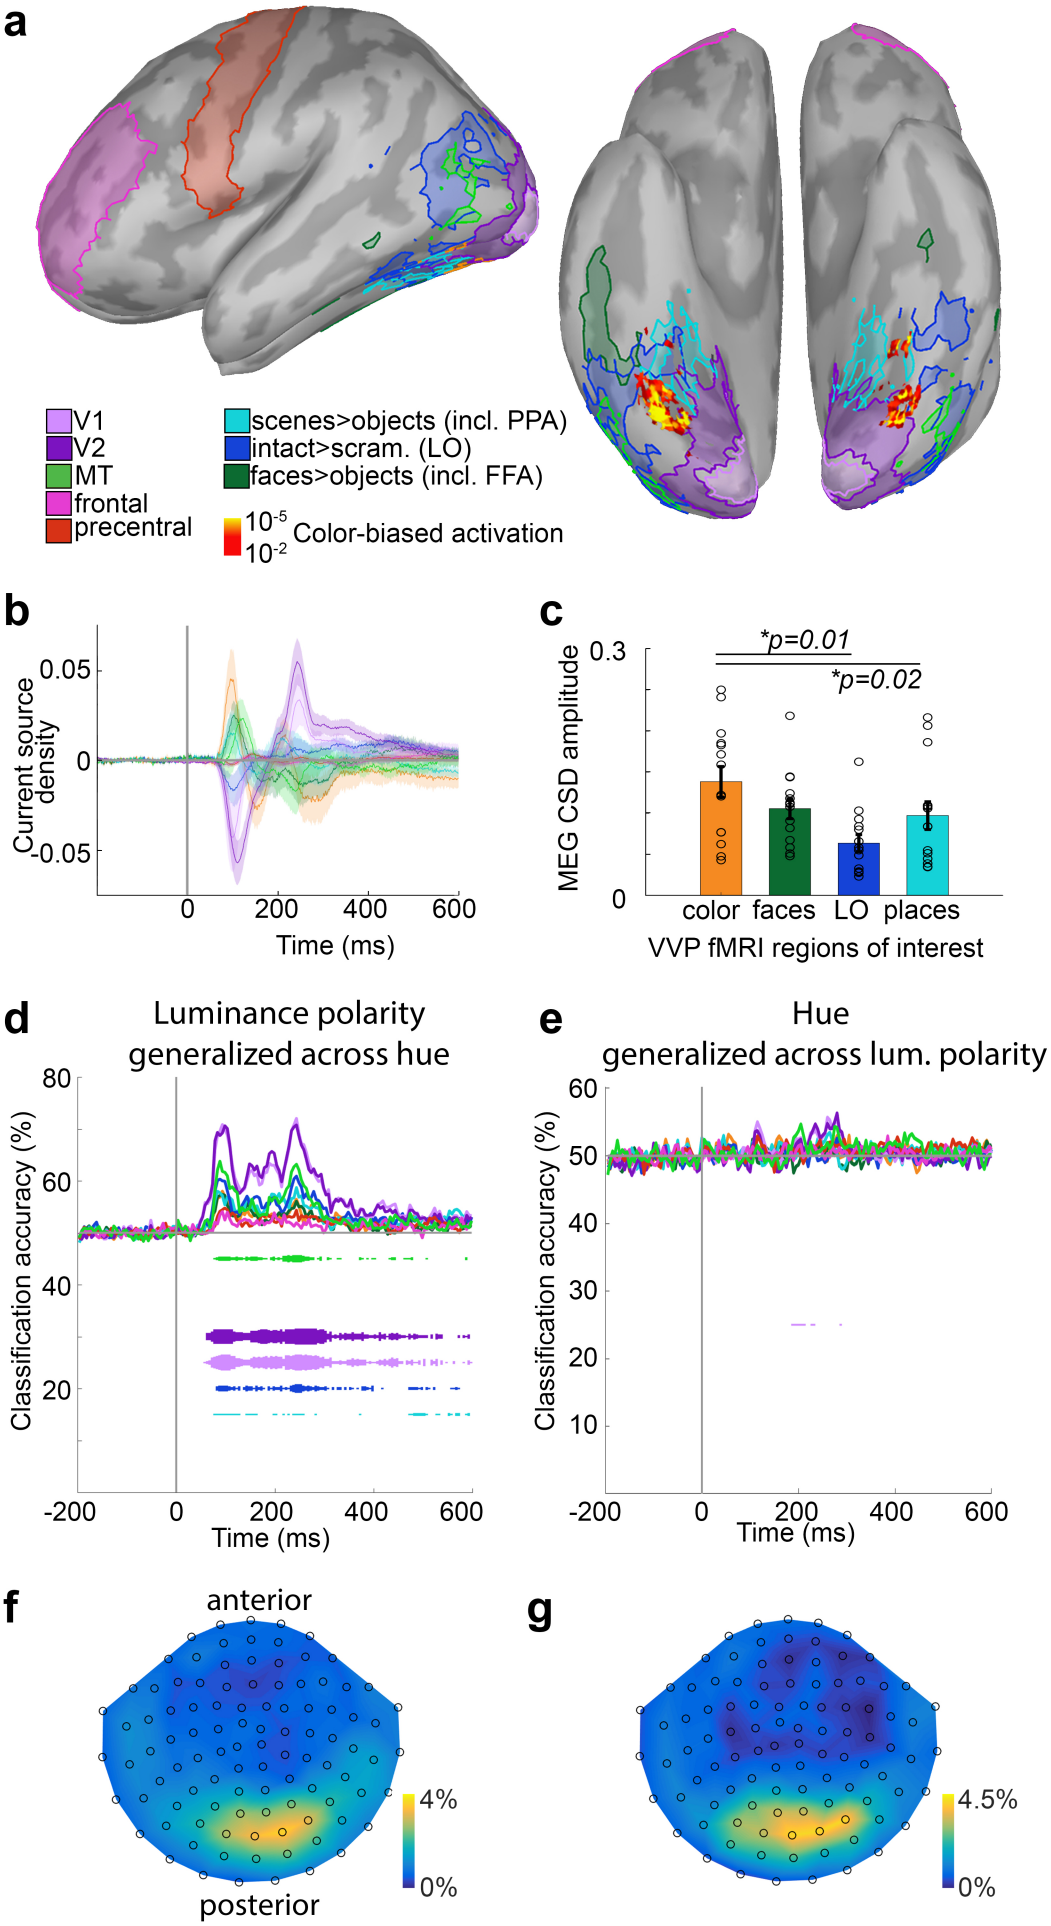
a
V1
V2
MT
frontal
precentral
scenes>objects (incl. PPA)
intact>scram. (LO)
faces>objects (incl. FFA)
10-5
10-2
Color-biased activation
b
Current source density
0.05
0
-0.05
0
200
400
600
Time (ms)
c
MEG CSD amplitude
0.3
0
*p=0.01
*p=0.02
color
faces
LO
places
VVP fMRI regions of interest
d
Luminance polarity
generalized across hue
Classification accuracy (%)
80
60
40
20
-200
0
200
400
600
Time (ms)
e
Hue
generalized across lum. polarity
Classification accuracy (%)
60
50
40
30
20
10
-200
0
200
400
600
Time (ms)
f
anterior
posterior
4%
0%
g
4.5%
0%

SUPPLEMENTARY TABLES and FIGURES

Temporal dynamics of the neural representation of hue and luminance polarity

Katherine L. Hermann[1*^], Shridhar R. Singh[1*], Isabelle A. Rosenthal[1*#], Dimitrios Pantazis[2], Bevil R. Conway[1,3, †]

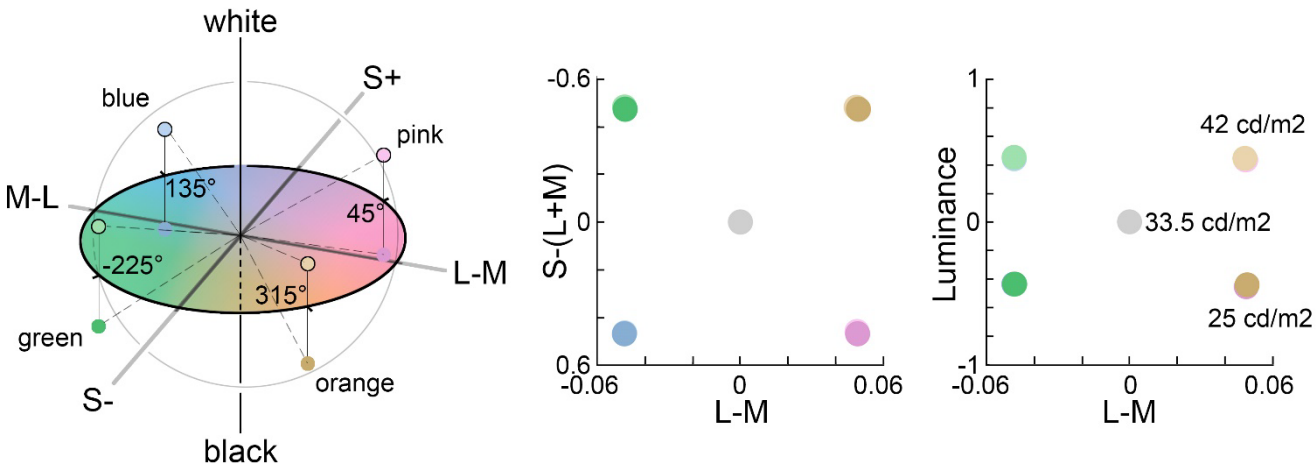

Supplementary Figure 1. Stimulus specification. Stimuli colors were set to the intermediate directions of the DKL color space, which is defined by the cardinal cone-opponent mechanisms.[1, 2] The left panel shows the eight colors at hue angles of 45°, 135°, 225°, and 315°, each hue at two luminance levels (elevations). The CIE xyY values for the stimuli are in SI Table 1. The DKL values were obtained using the function *lms2dkl* from Westland,[3] with the input LMS values computed in Supplementary Table 2. The right panels show the cone and luminance contrast of the stimuli, relative to the gray adapting background, with the luminance values in cd/m^2 indicated.

	CIE x	CIE y	CIE Y'	Y
Light pink	0.305229324	0.293448787	0.628205196	41.8950
Dark pink	0.304095045	0.264040315	0.370334398	24.6976
Light blue	0.272115611	0.307026806	0.630601697	42.0548
Dark blue	0.25265696	0.284531124	0.375699736	25.0554
Light green	0.310023683	0.426308239	0.63358717	42.2539
Dark green	0.312628272	0.502281154	0.375716373	25.0565
Light orange	0.352384707	0.400239965	0.631544192	42.1177
Dark orange	0.392438648	0.444792187	0.374272604	24.9602
Adapting gray	0.306634746	0.346610046	0.502314308	33.4993

Supplementary Table 1. The xyY values for the stimuli. The values in the third column (Y') yield luminance (Y), when multiplied by the maximum gun output (66.69 cd/2).

	L	M	S
Light pink	27.6333	14.2617	0.9216
Dark pink	16.5006	8.1970	0.6497
Light blue	26.7305	15.3243	0.9271
Dark blue	15.7224	9.3330	0.6554
Light green	26.8587	15.3952	0.4203
Dark green	15.7255	9.3310	0.1485
Light orange	27.7748	14.3429	0.4187
Dark orange	16.6733	8.2869	0.1469
Adapting gray	21.6919	11.8074	0.5389

Supplementary Table 2. The LMS coordinates of the stimuli, computed using the MATLAB function *xy2MB* from Westland.[3]

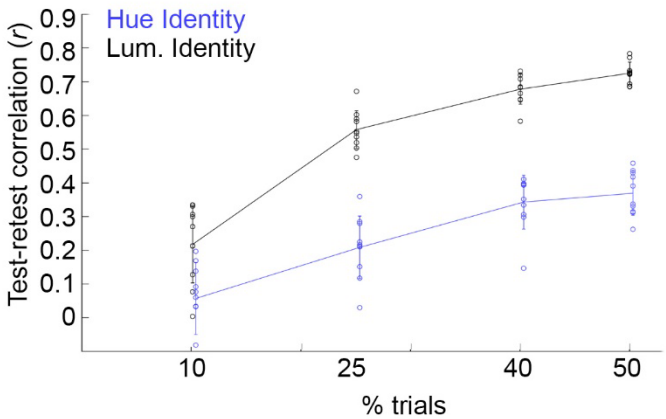

Supplementary Figure 2. Data reliability analysis for the pilot experiment (n=4 participants). The graph shows the correlation between pairs of independent data sets, as a function of the amount of data in the data set, for decoding hue identity and luminance-polarity identity. For each problem, pairs of classifiers were trained and tested on independent samples containing 10%, 25%, 40%, and 50% of the data, and the correlation between the classifiers' performance at each time point was calculated (i.e., the analysis shows the extent to which the shape of the classification curve is similar for independent data sets of different sizes). This procedure was repeated ten times to obtain the test-retest correlation and error bars, which represent the standard deviation across the ten replicates. This analysis provides a measure of data reliability.[4] The y-axis shows the correlation coefficient comparing the decoding magnitude at each point in the decoding time course between the pairs of data sets. The graph shows: (1) that the data are reliable (the test-retest estimates are above chance); (2) that the experiments have sufficient power to extract close to as much signal as is possible given the experimental conditions (the test-retest curves come close to plateau when extrapolated to 100% of the data); and (3) that to decode color, one needs relatively many trials. As in the main experiment, the stimuli were spirals, participants viewed each color 500 times, and each stimulus appeared for 100 ms followed by a 1 s ISI. Other details of experimental paradigm were the same as for the main experiment, with the following exceptions: the data were collected in 4 participants (3 female; separate people from those who participated in the main experiments); and the stimuli only included four colors of either a luminance increment or luminance decrement. The two hues were DKL angles 150 and 300, which are intermediate colors corresponding roughly to blue and yellow; the luminance of the stimuli were: background gray, 41 cd/m^2; positive luminance-contrast stimuli, 48-50 cd/m^2; and negative luminance-contrast stimuli, 30-32 cd/m^2).

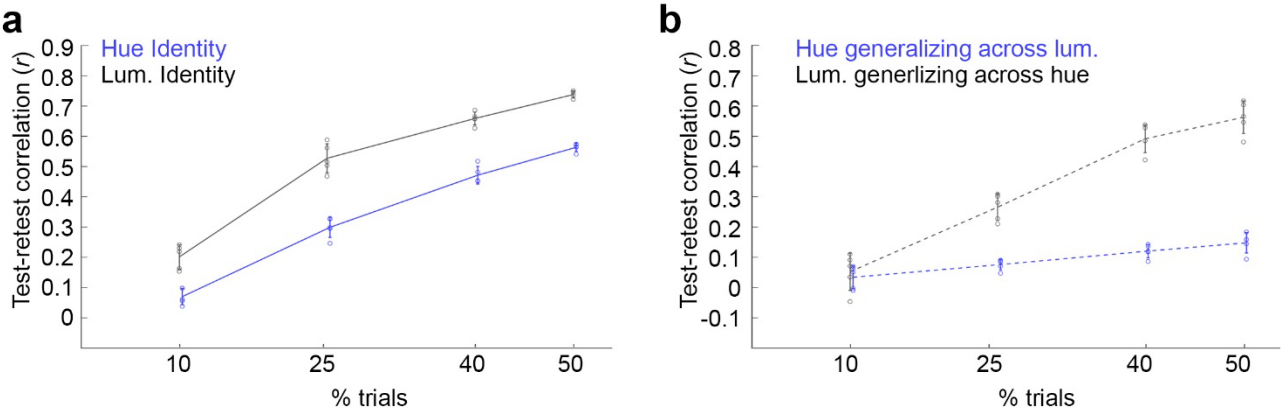

Supplementary Figure 3. Data reliability analysis for main experiments. **a)** Results for the identity problems. For each identity problem, pairs of classifiers were trained and tested on independent samples containing 10%, 25%, 40%, and 50% of the data, and the correlation between the classifiers' performance at each time point was calculated (the analysis shows the extent to which the shape of the classification curve is similar for independent data sets of different sizes). This procedure was repeated five times to obtain the test-retest correlation and error bars, which represent the standard deviation across the five replicates. **b)** Results for the generalization problems. Again, the correlation procedure was repeated five times to obtain the test-retest correlation and standard deviation error bars. The test-retest curves in panel (a) and (b) differ somewhat from those in Supplementary Figure 2, which shows the same analysis for the pilot experiment. The data for the main experiments were obtained in n=18 participants using 8 colors, whereas the data for the pilot experiment were obtained in four participants using 4 colors.

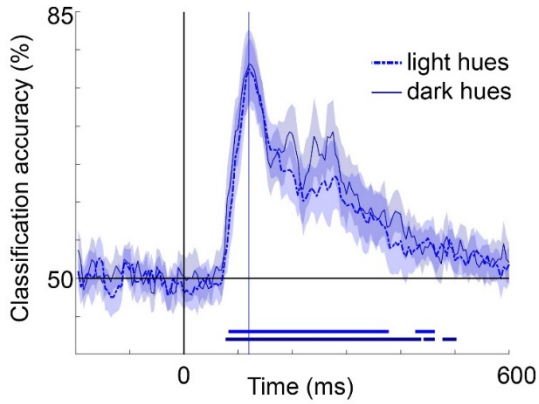

Supplementary Figure 4. The average classification accuracy for the 6 sets of hue identity problems between stimuli of the same luminance polarity (six problems among light hues; six problems among dark hues). Shading represents the 95% confidence intervals on the accuracy, which were obtained by bootstrapping 18 subjects 1000 times. The time to peak was not significantly different for identity problems among luminance-contrast increments (122 ms [115, 130]) compared to identity problems among luminance-contrast decrements (123 ms [115, 130]); and the magnitude of peak decoding accuracy for these two sets of problems also were not significantly different (for light stimuli: 77% [73, 82]; for dark stimuli: 78% [74, 83]). The horizontal sequence of data points above the x-axis shows time bins at which decoding was significantly above chance, FDR corrected, for a minimum of four consecutive time bins.

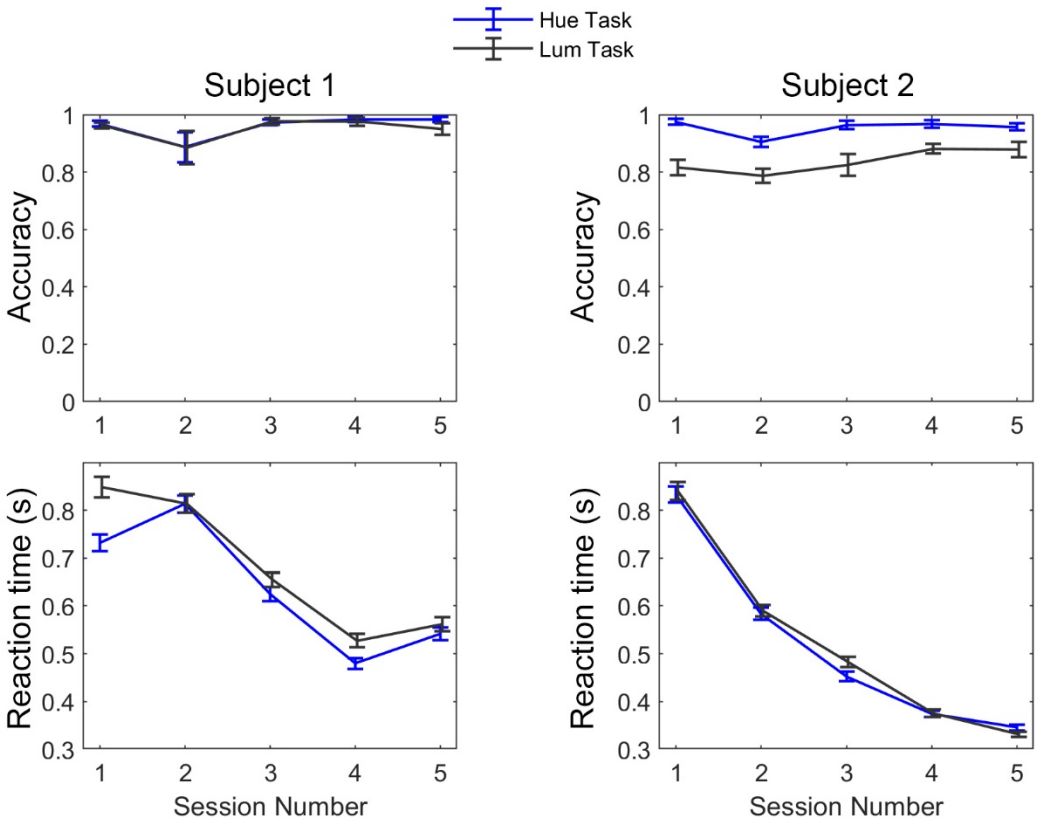

Supplementary Figure 5. Behavioral performance for the two subjects from the control experiment in which MEG data were collected while participants performed either the 1-back luminance-polarity matching task (black traces) or the 1-back hue matching task (blue traces; see Figure 5 for the MEG decoding results of this control experiment). Note the slightly longer reaction times for the luminance task (Subject 1) and the slightly lower accuracy for the luminance task (Subject 2), which reflects the general difficulty people face in making heterochromatic luminance-contrast matches. For subject 1, error bars and averages were obtained from n=249, 253, 248, 248, 251 trials for each session respectively. For subject 2, the number of trials per session were n=248, 254, 248, 248, 252. Error bars represent the standard error across trials.

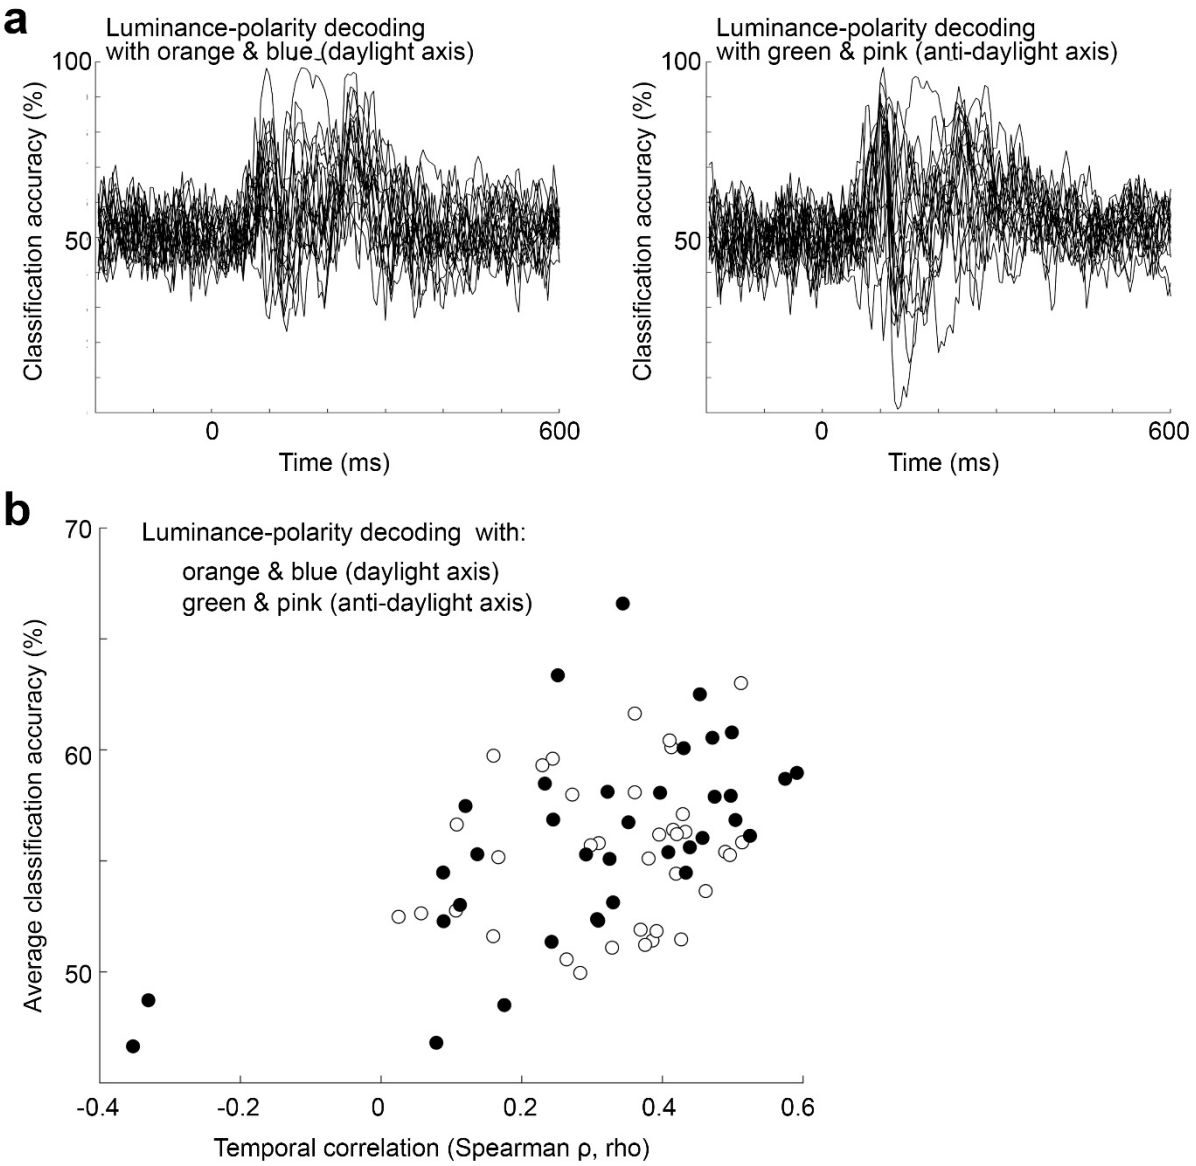

Supplementary Figure 6. Luminance-polarity decoding was more variable across participants for problems using orange and blue versus problems using green and pink. **a)** Individual luminance-polarity decoding curves for problems involving orange and blue (left panel) and pink and green (right panel) across the 18 participants (36 problems illustrated in each panel, two per participant involving orange and blue; and two per participant involving pink and green). The decoding curves using pink and green were more correlated across time, as quantified in panel (b). **b)** The average classification accuracy (averaged over the entire decoding time course from 0 to 600 ms) for each problem, versus the temporal correlation of the problem. The temporal correlation for each problem was computed by measuring the correlation between it and a problem drawn at random from the remaining 71 problems, repeated 1000x, and averaged. Decoding luminance polarity was more variable for data obtained using orange and blue compared to data obtained with green and pink (the black dots are shifted to the right; one-way MANOVA, p=0.0007).

Supplementary References

1. MacLeod DI, Boynton RM. Chromaticity diagram showing cone excitation by stimuli of equal luminance. *J Opt Soc Am* **69**, 1183-1186 (1979).

2. Derrington AM, Krauskopf J, Lennie P. Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology* **357**, 241-265 (1984).

3. Westland S, Ripamonti C, Cheung V. Computational colour science using matlab (2nd Edition). *John Wiley and Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, England*, (2012).

4. Norman-Haignere SV, Kanwisher N, McDermott JH, Conway BR. Divergence in the functional organization of human and macaque auditory cortex revealed by fMRI responses to harmonic tones. *Nature Neuroscience* **22**, 1057-+ (2019).