Combination of subcortical color channels in human visual cortex

Erin Goddard
School of Psychology, The University of Sydney, and the Australian Centre of Excellence in Vision Science, Sydney, Australia

Damien J. Mannion
School of Psychology, The University of Sydney, and the Australian Centre of Excellence in Vision Science, Sydney, Australia

J. Scott McDonald
School of Psychology, The University of Sydney, Sydney, Australia

Samuel G. Solomon
School of Medical Sciences and Bosch Institute, The University of Sydney, and the Australian Centre of Excellence in Vision Science, Sydney, Australia

Colin W. G. Clifford
School of Psychology, The University of Sydney, and the Australian Centre of Excellence in Vision Science, Sydney, Australia

Mechanisms of color vision in cortex have not been as well characterized as those in sub-cortical areas, particularly in humans. We used fMRI in conjunction with univariate and multivariate (pattern) analysis to test for the initial transformation of sub-cortical inputs by human visual cortex. Subjects viewed each of two patterns modulating in color between orange-cyan or lime-magenta. We tested for higher order cortical representations of color capable of discriminating these stimuli, which were designed so that they could not be distinguished by the postulated L–M and S–(L + M) sub-cortical opponent channels. We found differences both in the average response and in the pattern of activity evoked by these two types of stimuli, across a range of early visual areas. This result implies that sub-cortical chromatic channels are recombined early in cortical processing to form novel representations of color. Our results also suggest a cortical bias for lime-magenta over orange-cyan stimuli, when they are matched for cone contrast and the response they would elicit in the L–M and S–(L + M) opponent channels.

Introduction

Our rich experience of color includes the ability to discriminate and identify a diverse range of combinations of hue, saturation and luminance, yet our perceptual experience is based on the activity of just three categories of cone photoreceptor and the transformation of these signals by sub-cortical and cortical areas. At the sub-cortical level, there exist chromatically opponent channels (L–M and S–(L + M)) that carry information in parallel to visual cortex via the parvocellular and koniocellular layers of the LGN (Derrington, Krauskopf, & Lennie, 1984). Cortical mechanisms of color vision are generally less well understood, although psychophysical adaptation experiments indicate the existence of higher-order color mechanisms in the human visual system, which receive input from both the opponent channels of sub-cortical areas (Krauskopf, Williams, Mandler, & Brown, 1986; Webster & Mollon, 1991; Zaidi & Shapiro, 1993). In macaque it has been demonstrated that there are cells as early as V1 which prefer chromatic directions away from the cardinal directions that isolate the L–M and S–(L + M) mechanisms (Conway, 2001; de Valois, Cottaris, Elfar, Mahon, & Wilson, 2000), implying a combination of information from the sub-cortical channels early in visual cortex. For cortical cells to have chromatic preference intermediate to the cardinal axes, there must be some combination of the

L–M and S–(L + M) channels. In a recent fMRI study, Brouwer and Heeger (2009) found that the principal components of the response in V1 are consistent with a response dominated by an opponent coding of color, as found in sub-cortical areas, but that by hV4 and VO it more closely resembles our perceptual color space. That is, V1 shows a differential response to variations in color but not a continuous representation of hue, while in higher areas colors of similar hue evoke similar responses. This finding does not rule out the possibility that the signals of the fundamental pathways (L–M and the S–(L + M)) are combined in the early visual areas, such as V1, a possibility we address specifically in this study.

We obtained high resolution functional images of the BOLD (blood-oxygen-level-dependent) response from subjects' occipital and parietal lobes while they viewed colored stimuli. Previous studies of cortical chromatic mechanisms in humans have used perceptually relevant hues (Brouwer & Heeger, 2009; Parkes, Marsman, Oxley, Goulermas, & Wuerger, 2009). Our stimuli were not chosen for their perceptual relevance, but were designed to be metameric to the hypothesized sub-cortical chromatic mechanisms. Specifically, the stimuli were designed so as to fulfill the following conditions: (1) to induce the same magnitude of activity in the L–M opponent channel; and (2) to induce the same magnitude of activity in the S–(L + M) opponent channel. This was achieved by combining a given L–M modulation with a given S-cone isolating modulation in each of two different phases. When −S was in phase with M then the stimulus appeared lime-magenta; when +S was in phase with M then it appeared orange-cyan. The lime-magenta and orange-cyan stimuli can only be distinguished by the BOLD signal if there are cells which receive a combination of inputs from the L–M and the S–(L + M) pathways.

Univariate and multivariate analyses tested whether the BOLD response within each cortical visual area depended upon the color of the stimulus. Univariate analyses show what information about the stimulus is detectable in the average activity across a region, while multivariate classifiers are capable of also learning differences in the pattern of activation between stimuli. Multivariate classification analysis (for a review, see Haynes and Rees, 2006) is a useful tool to test for differences in the BOLD response of a visual area even where the mean activity of the area is not significantly different between stimuli, and has been used to infer the selectivity of different early visual areas for a range of basic visual attributes and their combination (Haynes & Rees, 2005a, 2005b; Kamitani & Tong, 2005, 2006; Mannion, McDonald, & Clifford, 2009; Parkes et al., 2009; Seymour, Clifford, Logothetis, & Bartels, 2009; Sumner, Anderson, Sylvester, Haynes, & Rees, 2008).

Both in the univariate and multivariate analyses employed here, an algorithm is trained to classify the stimulus from activity across a region, and tested on novel data. Above chance performance indicates that the area

contains information about the stimulus dimension that was varied. Here, our premise is that if we can use the activity across a visual area to discriminate between our stimuli then that area contains a representation of color that could only be generated through a transformation of the signals from the sub-cortical L–M and S–(L + M) pathways.

Materials and methods

Color calibration procedures and display system

Stimuli were generated and displayed using Matlab (version 7) software on a Dell Latitude laptop computer driving an nVidia Quadro NVS 110M graphics card to draw stimuli to a 35 × 26 cm Philips LCD monitor, with 60 Hz refresh rate, viewed from a distance of approximately 1.58 m. Scanning took place in a darkened room. Subjects, while lying in the scanner, viewed the monitor through a mirror mounted above the head cage which reflected the image from the LCD monitor located behind the scanner. Stimuli were calibrated in situ for the LCD monitor and mirror arrangement, using measurements obtained with a PR-655 SpectraScan spectrophotometer (by Photo Research Inc.).

Changes in both chromaticity and luminance of the screen with increasing R, G and B values were taken into account when generating the experimental stimuli. The CIE (xyY) coordinates measured for 16 values during calibration were interpolated to 255 values using the best-fitting spline, and these were used to calculate the luminance and chromaticity for each combination of R, G and B intensity values.

Data were collected on five subjects (three male), aged between 24 and 40 years, with normal or corrected to normal visual acuity and normal color vision, as tested using Ishihara plates (Ishihara, 1990). All subjects provided informed consent, and the entire study was carried out in accordance with guidelines of the University of Sydney Human Research Ethics Committee.

Chromatic, spatial and temporal stimulus properties

Example stimuli are shown in Figure 1. The stimulus was an annulus, centered on the fixation point, with an inner diameter of 0.8 degrees visual angle, and an outer diameter of 7.8 degrees. The remainder of the screen was held at the mean luminance, which was 6.78 cd/m^2 [CIE (1931) *x, y*∼0.300, 0.337], and all stimuli were produced by spatiotemporal modulation around this point. The annulus contained a spatial pattern that counterphased
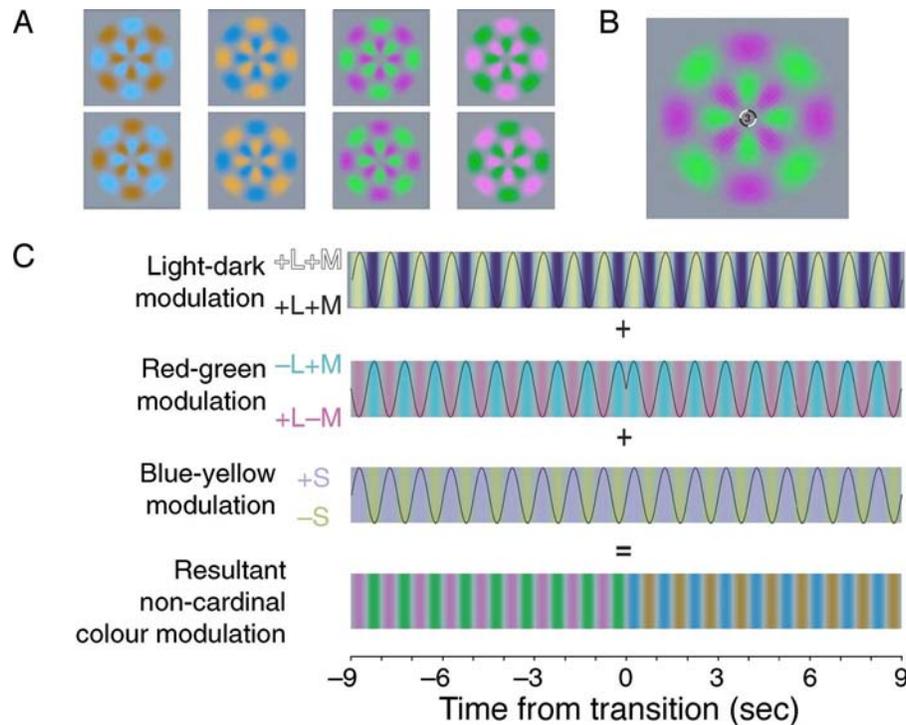
Figure 1. Stimuli used in the fMRI experiment; **A**: the color of the stimulus modulated sinusoidally between the upper plaid and the lower plaid at 1 Hz. The stimuli on the left are orange and cyan, on the right, the stimuli are lime and magenta. For both color pairs, minimum motion was used to determine each subjects' perceived equiluminance point, and a 25% luminance modulation was added. In the first and third pair of stimuli the light/dark modulation is paired with cyan/orange and lime/magenta, respectively. In the second and fourth stimuli these pairings are reversed. **B**: Example stimulus with fixation task. At fixation, there was a light gray cross surrounded by a high contrast ring, as illustrated above. The high contrast ring provided feedback to subjects when they made small eye movements, since an afterimage would become visible. While subjects fixated on the central cross (partially obscured by the digit), they were required to respond with a button press whenever either of two target items were presented in the digit stream. Digits updated at 3 Hz, were presented in random order, and could be 0 to 9 inclusive, each in either black or white. The target items were a conjunction of a digit and a particular color, for example either a black 3 (shown here) or a white 7. The fixation task was unrelated to the experimental stimulus, which was presented in the annular region surrounding fixation. **C**: Modulation of cardinal sensitive mechanisms over time in the experimental stimulus, for an example 18 seconds including a transition from a light magenta-dark lime block to a light cyan-dark orange block. At each transition, the non-cardinal modulation switched from a lime-magenta block to an orange-cyan block or vice versa. At each transition there was a phase reversal in either the L–M or the S–(L + M) cardinal modulation; here there is a phase reversal in the L–M modulation at the time of transition. The luminance (light–dark) modulation had no reversals at any time. The amplitude of response of any cardinal sensitive mechanisms should be constant across the lime-magenta and orange-cyan blocks. Phase reversals are the only cue that could be used to discriminate the stimuli where the cardinal sensitive mechanisms are kept independent, but the block order was balanced such that this cue could not be used to predict the stimulus. We used one of two block orders for each run, where the second was a reversal of the first order.

sinusoidally at a temporal frequency of 1 Hz. The spatial pattern was the multiplication of a radial and a concentric sinusoidal modulation, (the resultant plaid pattern is shown in Figures 1A and 1B). All these modulations can be represented in a three-dimensional color space described previously (Derrington et al., 1984; DKL space). Along the L–M axis only the signals from L and M cones vary, in opposition, without variation in luminance. Along the orthogonal S-cone isolating axis there is no modulation of either the L or M cones. The L–M and S axes define a plane in which only chromaticity varies. Normal to this plane is the luminance axis along which the signals from

the L and M cones vary in proportion. The axes were derived from the Stockman and Sharpe (2000) 2-degree cone spectral sensitivities and adjusted individually for each observer (see below). The scaling of these axes is largely arbitrary; here we used modulations along the isoluminant axes that were 90% of the maximum modulation achievable within the gamut of the monitor. Modulation along the L–M axis produced maximum cone contrasts of 15.4% in the L-cone and 17.8% in the M-cone; along the S-cone axis the maximum S-cone contrast was 79.6%. Cone contrast values for all stimuli are listed in Table 1. Each frame of the stimulus was generated prior

t1.1	Stimulus Color	L-cone Contrast	M-cone Contrast	S-cone Contrast
t1.2	Dark Cyan	−0.154	0.015	0.688
t1.3	Light Cyan	0.031	0.178	0.796
t1.4	Dark Orange	−0.031	−0.178	−0.796
t1.5	Light Orange	0.154	−0.015	−0.688
t1.6	Dark Magenta	−0.031	−0.178	0.688
t1.7	Light Magenta	0.154	−0.015	0.796
t1.8	Dark Lime	−0.154	0.015	−0.796
t1.9	Light Lime	0.031	0.178	−0.688

t1.10　Table 1. Cone contrast values for stimuli calibrated for the subjective equiluminance point of observer EG. The background of the stimuli had CIE xy coordinates of 0.30, 0.34, and a luminance (Y) of 6.78 cd/m^2.

to the experiment as a bitmapped image, and then these images were drawn to the screen for each stimulus presentation using routines from PsychToolbox 3.0.8 (Brainard, 1997; Pelli, 1997).

There were four stimuli, chosen such that over time all four would: (1) equally stimulate the L–M opponent channel; and (2) equally stimulate the S–(L + M) opponent channel. The angle of each stimulus within the isoluminant plane was intermediate to that of the L–M and S–(L + M) axes, and was defined as a vector addition of modulations along those two axes. When the L–M modulation was in phase with the S–(L + M) modulation the stimulus modulated between magenta and lime; when these pairings were reversed the modulation appeared orange-cyan. The point of subjective isoluminance (the angle of the isoluminant plane from the luminance axis) was estimated separately for each observer using the minimum motion technique described by Anstis and Cavanagh (1983), for the magenta-lime and for the orange-cyan modulations. A 25% luminance modulation was then added to the subjectively defined isoluminant modulation in one of two phases. The four stimuli therefore appeared: light magenta–dark lime, dark magenta–light lime, light orange–dark cyan, and dark orange–light cyan (example stimuli in Figure 1A, and example cone contrast values in Table 1).

The luminance modulation was added so that if there were any residual differences in luminance between the lime-magenta and orange-cyan blocks, this difference should be masked by the much larger luminance modulation. The contrast response of early visual areas to luminance defined stimuli is steeper at low than high contrast (for example, see Liu & Wandell, 2005). Thus a luminance artifact in our stimuli would result in a much smaller difference in the response to the two stimuli than if there was no luminance modulation (and hence lower luminance contrast). The same rationale underlies the use of random luminance noise to mask potential luminance artifacts (Birch, Barbur, & Harlow, 1992; Kingdom, Moulden, & Collyer, 1992; Sumner et al., 2008). For the

analyses shown in Figure 4 the classifier was trained and tested with two groups of blocks: one group included the two types of lime-magenta blocks (the two types differed in the relative phase of the luminance modulation) and the other group included the two types of orange-cyan blocks. We also performed a control analysis, where the classifier was trained to discriminate lime-magenta vs. orange-cyan on blocks that had only one luminance phase (one of the two types of lime-magenta blocks, and one of two types of orange-cyan blocks), and then tested on its ability to discriminate the other two types. The results of this analysis are shown in Appendix B.

Experimental design

All experimental scans were completed during a single session for each subject. The session included ten functional scans, each lasting 4.5 minutes. During each scan the subject viewed 18 blocks of the experimental stimulus, alternating between orange-cyan and lime-magenta blocks. Each block was 15 seconds long, and data from the first and last block were excluded from our analysis. In order to change the color of the stimuli, either the L–M or the S–(L + M) modulation must change phase, as illustrated in Figure 1C. This phase reversal is likely to induce an increased response of a neural population which responds to the relevant cardinal color, as in the characteristic response rebound used in fMRI adaptation (for example, see Engel & Furmanski, 2001; Kourtzi, Erb, Grodd, & Bülthoff, 2003). It is also possible that the response to the phase reversal may be evident in the BOLD signal, even when averaging across activity within a block, and could potentially be used to discriminate the two types of stimuli (for example, if an orange-cyan block always commenced with a L–M phase reversal). In order to eliminate this potential source of information about the color of the stimuli, we balanced the stimuli so that the pattern of phase reversals could not be used to predict the stimulus. We used one of two block orders for each run, where the second was a reversal of the first order.

Localizer

To select those voxels in each visual area most responsive to the experimental stimuli, two additional localizer scans during the same session as the experimental scans for each subject. Using a localizer scan that is separate from the experimental data avoids circularity that could otherwise be present (Kriegeskorte, Simmons, Bellgowan, & Baker, 2009). Localizer scans included a total of 17 blocks of 15 seconds each, comprised of stimulus blocks interleaved with blocks of fixation only. The stimulus blocks included 2 lime-magenta blocks and 2 orange-cyan blocks, in addition to 4 black-white blocks

where the stimulus had the same spatial arrangement but was modulated between black and white.

Fixation task

Throughout experimental and localizer scans, subjects performed a task at fixation that was unrelated to the annular experimental stimulus. This task was designed to be attentionally demanding in order to direct attention away from the experimental stimulus as much as possible. While this would have reduced the BOLD response, and so likely decreased the ratio of signal to noise in the data which were input to the classifier, it greatly reduces the chance that subjects could have systematically directed more attention to one type of stimulus block.

Subjects were required to detect a conjunction of contrast polarity and number in a digit stream of the digits 0 to 9 inclusive, presented at fixation, updated at 3 Hz. Digits were either black or white, against the mean gray of the background, as seen in Figure 1B, and the order was randomly generated for each run. Subjects responded with a button press to the onset of either of two target digits, one only when black and the other only when white (for example a black 3 or a white 7). Responding to a conjunction of digit and contrast polarity made this a difficult task. Target digits were updated at the beginning of each run to increase task difficulty and minimize practice effects.

All subjects performed the task significantly above chance ($p < 0.01$, permutation test), demonstrating that they were engaged in the task, but each subject also made errors, implying that the task was not trivial and required attention. For no subject was there a significant difference in performance between lime-magenta and orange-cyan blocks, consistent with equal attentional resources being devoted to the task in each case.

fMRI methods

fMRI data were collected using a 3T Philips scanner (Symbion Imaging Centre, Prince of Wales Medical Research Institute, Sydney, Australia), with a birdcage head coil.

Anatomical measurements and definition of gray matter

The anatomical image for each subject was generated from the average of three scans. Two of these were high resolution ($1 \times 1 \times 1$ mm) structural MR images of each subject's whole brain, acquired using a Turbo Field Echo (TFE) protocol for enhanced gray–white contrast. A third, higher resolution ($0.75 \times 0.75 \times 0.75$ mm) scan of the caudal half of the head was also acquired in order to recover more anatomical detail of the occipital lobes.

Using the Statistical Parametric Mapping (SPM) software package SPM5 (Frackowiak, Friston, Frith, Dolan, & Mazziotta, 1997), anatomical images were each reoriented to approximately the same space using anterior and posterior commissures as anatomical landmarks. Fine alignment of these anatomical images was carried out using normalized mutual information based coregistration, and each of the anatomical images were resampled so that they were in the same voxel space with a resolution of $0.75 \times 0.75 \times 0.75$ mm. From each image we removed intensity inhomogeneities using a nonparametric inhomogeneity correction method (Manjón et al., 2007), and normalized the images such that the white matter had an approximate intensity of 1. The coregistered, inhomogeneity corrected, normalized images were then averaged together to produce a mean anatomical image for each subject.

ITKGray software (Yushkevich et al., 2006) was used to define the white matter of each subject, initially using automatic segmentation tools, then using manual editing. The segmentation image was imported into mrGray, part of the mrVista software package developed by the Stanford Vision and Imaging Lab (http://white.stanford.edu/software/). In mrGray, gray matter was 'grown' out from the white matter in a sheet with a maximum thickness of 4 voxels.

Functional measurements

fMRI data were acquired using a T2*-sensitive, FEEPI pulse sequence, with echo time (TE) of 32 ms; time to repetition (TR) of 3000 ms; flip angle 90; field of view 192 mm $\times$ 70.5 mm $\times$ 192 mm; effective in-plane resolution 1.5 mm $\times$ 1.5 mm, and slice thickness 1.5 mm. 47 slices were collected in an interleaved, ascending order, in a coronal plane tilted such that the scan covered the whole of the occipital lobe and the posterior part of the parietal and temporal lobes. Using SPM5, all functional data were preprocessed to correct for slice time and head motion before alignment to the structural data. Data from functional scans were aligned to a whole head anatomical scan acquired in the same session, using normalized mutual information based coregistration. The functional data were then aligned to the subject's average anatomical by first aligning the within session anatomical with the average anatomical scan, then applying the same transformation to the functional data.

Definition of retinotopic areas

For each subject, the precise anatomical locations of the early areas of visual cortex (V1, V2, V3, V3A/B, hV4, and VO) were found functionally using standard retinotopic mapping procedures (Engel, Glover, & Wandell, 1997; Larsson & Heeger, 2006; see Wandell, Dumoulin, & Brewer, 2007, for a summary). Subjects were scanned while viewing first a rotating wedge then an expanding ring stimulus, overlaid on a fixation cross of light gray lines, as shown on the key above the maps in Figure 2 (Schira, Tyler, Breakspear, & Spehar, 2009).
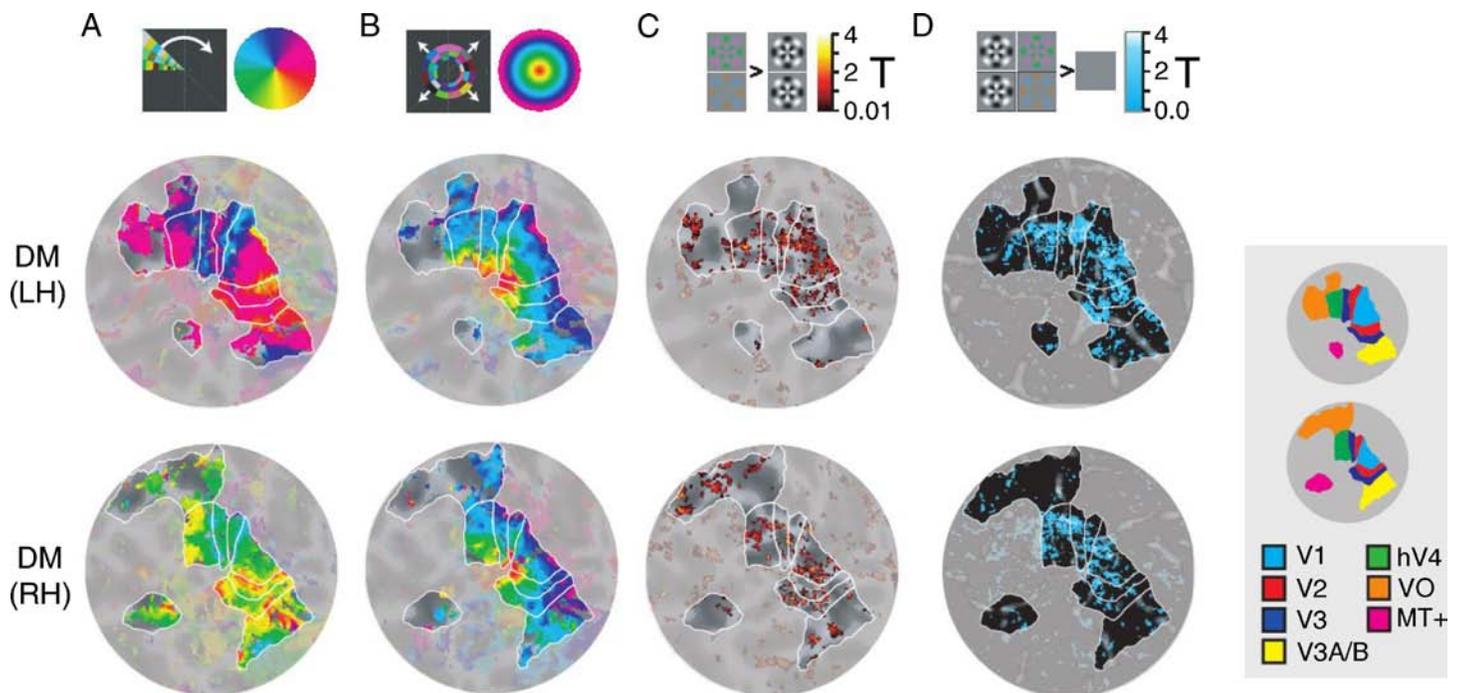
Figure 2. Example maps of functionally defined retinotopic areas for the left and right hemispheres of subject DM. In each of **A**, **B**, **C** and **D** the underlying grayscale image shows the flattened map of visual cortex, centered on the occipital pole; the darker the gray the deeper the sulcus. In **D** the grayscale anatomical map was darkened to increase the visibility of the overlaid image. **A** & **B**: Flattened maps of visual cortex overlaid with phase maps of the response to the wedge and ring stimuli, respectively. Above these maps is a schematic of the stimulus (top left) and a color map showing the area of the visual field which each color corresponds to in the phase maps (top right). In **C** the same flattened map of visual cortex is overlaid with a heat map showing those voxels which responded more to the chromatic than the achromatic stimuli; the significance of this result for each voxel is indicated by the T-statistic color map above. Areas V1, V2, V3, V3A/B and hV4 were defined on the basis of the wedge and ring phase maps in **A** and **B**, while area VO was defined according to a combination of the wedge and ring phase maps in **A** and **B**, and the contrast in **C**. MT+ was defined according to a motion versus static dots localizer (not shown). The borders of each of these areas are drawn on each of the maps in **A**–**D**; the key on the right indicates which of the outlined regions corresponds to each visual area. In **D**, the heat map indicates those voxels that were included in the analysis: those which responded significantly more to chromatic or achromatic versions of our stimuli than to fixation; the significance of this result for each voxel is indicated by the T-statistic color map above.

Averaged data from the wedge and ring stimuli were smoothed with a Gaussian kernel of half width 1.5 mm, then projected onto a computationally flattened representation of the cortex for each hemisphere of each subject, using mrVista. Areas V1, V2, V3, V3A/B and hV4 were manually defined on the phase and eccentricity maps derived from the wedge and ring stimuli (shown for an example subject in Figures 2A and 2B, respectively), using the conventions described by Larsson and Heeger (2006). According to these definitions the foveal representation at the occipital pole is shared by V1, V2, V3, and hV4, while V3A and V3B, which border the dorsal part of V3, share a dorsal fovea. For our analysis we did not attempt to separate V3A and V3B. We defined hV4 as a ventral hemifield representation that borders the ventral part of V3.

For area VO, the phase and eccentricity maps were considered in conjunction with a flattened map of those voxels that responded more to chromatic than achromatic stimuli in the localizer scan (as shown for an example subject in Figure 2C). Where it existed, we used the hemifield representation from the phase and eccentricity maps to define VO. Where the retinotopic map from the wedge and ring stimuli was unclear, we tended towards a liberal definition of VO, in order to avoid excluding any voxels in the region which showed a preference for chromatic stimuli in the localizer analysis. Each retinotopic area was defined on the flattened map of a subject's cortex then transformed into the space of the subject's anatomical, smoothed (FWHM = 1.5 mm), and resliced to the resolution of the functional images using 4th degree B-spline interpolation. Voxels assigned to each visual area were allocated a value reflecting the cumulative influence of such transformations. To prevent overlapping voxels between adjacent visual areas, each voxel was assigned to the visual area for which it possessed the greatest value.

Area MT+ was defined on the basis of a separate localizer scan in which blocks of low contrast static and moving dots were interleaved with fixation only blocks. In SPM5, we specified a general linear model of this data, and defined MT+ by finding lateral clusters of voxels that responded

more to moving than to static dots. The definition of MT+ was projected onto the flattened map for the purposes of illustration, as in Figure 2, but the original 3D definition of MT+ was used in the analysis of the functional data.

Functional preprocessing

Data from each of the two localizer scans were processed using the methods described above, then analyzed using a General Linear Model (GLM) in SPM5. We pooled responses to the luminance and chromatically defined stimuli and contrasted these with the response to the fixation only blocks. The subsequent analyses included only those voxels that responded significantly more ($p < 0.05$, uncorrected) to the stimulus than fixation only. These voxels are shown on an example flattened map for one subject in Figure 2D.

For all functional scans the BOLD signal was labeled with the stimulus presented 2 images (6 seconds) previously, in order to compensate for the delayed hemodynamic response, then was highpass filtered (cutoff 128 s, using filtering methods from SPM5) in order to remove low frequency confounds in the data, and finally converted into z-scores for each of the ten runs in order to reduce variability from inter-run differences. Data from each voxel were z-scored separately. Within each 15 second block the BOLD response, normalized according to the procedures described above, was averaged across the 5 measurements to give a single score for each block in each run.

Classifier analysis

Classifiers were restricted to each of several functionally defined visual areas for each subject and trained to discriminate the two patterns. We compared the performance of classifiers trained on two types of data for each area: in the univariate case, the classifier was trained and tested on the average activation across voxels within an area (that is, 1 value per block), while in the multivariate case the classifier was trained and tested on the pattern of activity across voxels within an area (n values per block).

We used a linear support vector machine (SVM) classification technique in our analysis. Support vector machines are powerful in their ability to learn a decision rule for multivariate data (Bennett & Campbell, 2000): in our case, for n voxels with 144 data points each (72 from lime-magenta and 72 from orange-cyan blocks) they learn the hyperplane which best separates the data points in an n dimensional space, where each dimension corresponds to the normalized response of one voxel (using linear SVMs, we require that the hyperplane's projection onto any two dimensions is linear). We evaluated classifier performance on its ability to generalize, i.e. to correctly discriminate data that were not included in the training set. For the univariate classifier the technique was the same, but there was only one dimension along which the 144 data points

varied, so the power of the support vector machine was not utilized.

Cross-validation: Leave-one-out train and test

Classification analysis was performed using a Matlab (version 7.5) interface to SVM-light 6.01 (Joachims, 1999). The classifier was trained on the scores from 9 runs and tested on a tenth; this procedure was repeated 10 times. This leave-one-out train and test procedure resulted in the data from each run being used as test data once, giving an average classifier performance (reported in the Results section as a percentage correct) based on the accuracy across 160 classifications, while ensuring that the test data never included data that were used in training.

We repeated the classification analysis with increasing numbers of voxels (n) within each visual area, from $n = 3$ voxels, to the $n = \max N$, where $\max N$ is the total number of voxels that reached significance in the localizer analysis. Voxels in each area were ranked according to their t statistic from the localizer analysis, based on the separate localizer scans, in order to select voxels that responded to the area of visual field occupied by our experimental stimuli and exclude those which represented areas of the visual field which were more foveal or peripheral than our annular stimuli. The top n most significant voxels were used in each case. Classifier performance generally increased as more voxels were included in the analysis, but there was some variability around this trend. To summarize classification performance (as reported below) we fit the classifier performance (P) as number of voxels (n) increased with an exponential growth function which reaches a limit (L), given by

$$P = 0.5 + (L - 0.5)(1 - \mathrm{e}^{-n/c}), \tag{1}$$

where 0.5 is chance performance (at $n = 0$), and c is a curvature parameter, specifying how many voxels the curve takes to reach the limit, L. When the curve fit the data, the classifier performance reported below corresponds to the limit (L) of the growth function. When the curve could not be fit to the data within 100 iterations of the Matlab function *nlinfit* (usually when classifier performance was low), the average classifier performance, rather than the limit of the curve, is reported as the summary statistic of classifier performance. Classifier performance as a function of number of voxels, along with the best-fitting curve, is plotted in Appendix A.

Permutation analysis

To test the statistical significance of classifier performance, we ran a permutation analysis to estimate classification accuracy expected by chance alone. Permutation tests are non-parametric and so do not include an assumption of normality, and such tests have previously been employed to evaluate classification analysis (Mourao-Miranda, Bokde,

Born, Hampel, & Stetter, 2005). For each area in each subject we performed the same analysis as that described above, except that before training the classifier we randomly permuted the stimulus labels associated with each block in the training data set. Using 1000 repetitions of this permutation analysis, we generated a population of 1000 estimates of the classifier accuracy that could be expected in cases where the data did not contain any stimulus-related information. For each iteration of the permutation analysis we averaged these estimates across subjects for each area and compared these 1000 values with the observed between-subject mean accuracy. In the statistics reported below for classifier performance, p-values were calculated by finding the proportion of these 1000 estimates which were greater than the observed classifier accuracy.

Results

We tested for the presence of cortical representations of color capable of discriminating stimuli that cannot be distinguished by the L–M and S–(L + M) sub-cortical opponent channels. Our lime-magenta and orange-cyan stimuli contain the same L–M and S–(L + M) modulations, varying only in the phase that these modulations are added together. For the lime-magenta stimuli, the L–M and S–(L + M) modulations are the same phase, whereas for the orange-cyan stimuli the L–M and S–(L + M) modulations were in opposite phase. We used fMRI to measure changes in BOLD signal as an indirect measure of neural activity, then asked the extent to which the visual stimulus could be discriminated from patterns of brain activity in a predefined visual area.

Below, we report stimulus related differences in both the mean activity and pattern of activity across a range of regions. There was a small but reliable bias across subjects for lime-magenta over orange-cyan stimuli in the mean activity across each region, and we found that the difference in the mean activity was sufficient for a univariate classifier to learn to correctly discriminate the stimuli. We also found evidence of additional stimulus-related information in the pattern of activity across V1 and V2, using multivariate classifiers.

Consistent bias for lime-magenta over orange-cyan stimuli in average activity

There was a significant difference in the response to lime-magenta vs. orange-cyan stimuli, impossible without combination of signals from the fundamental cone-opponent channels. For the univariate analysis we averaged across those voxels within each area for which there was a significant difference in their response to the localizer stimulus versus fixation. The average z-scored activity across an area for each block was treated as a separate measure of the area's response to lime-magenta or orange-cyan stimuli,

giving 80 measurements for each; the distributions of these measurements for each subject are shown in Figure 3.

In all subjects each area that showed a significant bias for one color modulation showed the same bias; signal was greater for lime-magenta than for orange-cyan. As shown in Figure 3, the mean of the 80 lime-magenta blocks was significantly greater than the 80 orange cyan-blocks ($p < 0.05$, two-tailed t-test); in 25 areas across 5 subjects. The only area for which there was not a significant difference between the two types of stimuli for any subject was area MT+. We conclude that stimuli which equally modulate the cardinal axes of color space are not equally represented in visual cortical areas, and this biased representation is seen as early as V1.

Consistent with the differences in the average response across each region, univariate classifiers performed significantly better than chance as early as V1, as shown in Figure 4 (light gray bars). The average univariate classification performance across subjects in all areas was significantly above chance ($p < 0.05$, one-tailed permutation test). MT+ showed the lowest performance for 4 of 5 subjects.

For 4 of 5 subjects, the area with the best classifier performance was V2. The earlier cortical visual areas (V1, V2 and V3) generally outperformed the dorsal area V3A/B, as well as the ventral areas hV4 and VO. V1, V2 and V3 generally had a greater number of voxels, which may account for their high performance. In order to test this, we repeated the classifier analysis on 100 voxels that were randomly chosen from those voxels in the area that responded to the localizer stimulus. For this subset of 100 voxels, the classifier accuracy averaged across subjects in V1, V2 and V3 (61, 64 and 61%) was still better than in V3A/B, hV4 and VO (45, 56 and 52%). This suggests that differences in classifier performance cannot be accounted for by the generally greater number of voxels included in the analysis for areas V1, V2 and V3. Area MT+ had fewer than 100 voxels that responded to the localizer in all subjects, and so was excluded from this reanalysis.

Additional information in the pattern of activity for areas V1 and V2

Multivariate pattern classifiers were also trained to discriminate the two types of stimulus, allowing us to test for additional stimulus related information in the pattern of activity across each area. The univariate classifier was trained on a subset (9 of 10 runs) of the average data (as plotted in Figure 3), and tested on the remainder. The multivariate classifier was trained and tested in the same way on data which was not averaged across voxels, so that in addition to the average it can learn differences in the pattern of activity across an area between blocks. This analysis allows a comparison of the results from the univariate and multivariate classification techniques, which gives an indication of how the information in the pattern of activity across an area differs from information given by the mean response.
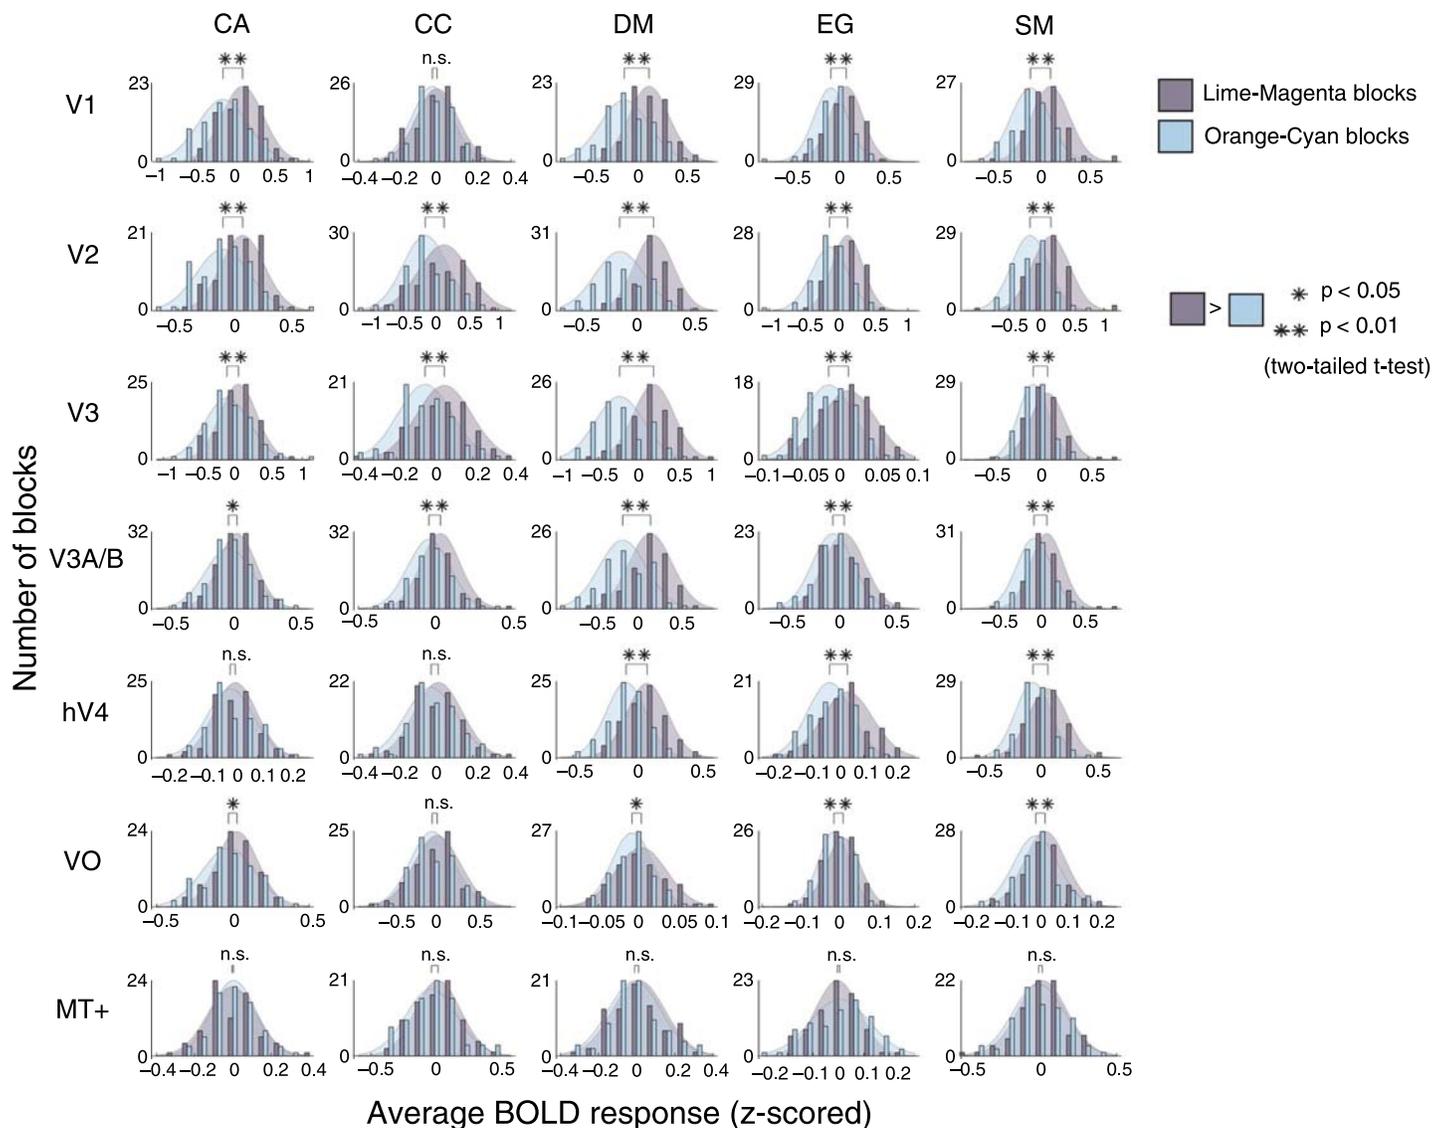
Figure 3. Mean activity in response to lime-magenta vs. orange-cyan blocks, across different cortical visual areas, for each subject. For each block, the z-scored BOLD response was averaged across all voxels for which there was a significant difference in their response to the localizer stimulus versus fixation. The histograms plot the distribution of these averages. Behind each histogram, a normal distribution of the same mean and standard deviation is plotted for reference. For all subjects, where there was a significant difference between the response to lime-magenta and orange-cyan blocks (tested with a two-tailed t test), the response to the lime-magenta blocks was greater than the response to orange-cyan blocks. The functional data did not include a fixation block; the average % signal change in the color blocks in the localizer data, relative to fixation (± 1 standard error of the between subject mean) were V1: 0.65 (± 0.16); V2: 0.78 (± 0.19); V3: 0.56 (± 0.17); V3A/B: 0.18 (± 0.16); hV4: 0.67 (± 0.27); VO: 0.35 (± 0.19); MT+: -0.09 (± 0.17).

Multivariate classification performance across areas (see Figure 4) followed a similar trend to that found for the univariate classifiers; earlier visual areas (V1, V2 and V3) tended to outperform V3A/B, hV4 and VO. Classification accuracy was poorest in MT+, where performance was not significantly different from chance. This trend was also found when the classifier was trained and tested on one hundred voxels, randomly chosen from those voxels in the area that responded to the localizer stimulus: the average between-subjects classifier accuracy in V1, V2 and V3 (63, 67 and 59%) was still better than in V3A/B, hV4 and

VO (57, 56 and 49%). Classification performance was generally higher for multivariate classifiers, and this difference was significant ($p < 0.01$, two-tailed permutation test) for areas V1 and V2. In VO and MT+, the univariate classifier outperformed the multivariate classifier, but this difference was not significant.

Since there was no requirement on our classifiers to predict an equal number of orange-cyan and lime-magenta blocks, it was possible for classification performance to be better for one type of test stimulus. For example, if classification of orange-cyan test stimuli were at chance
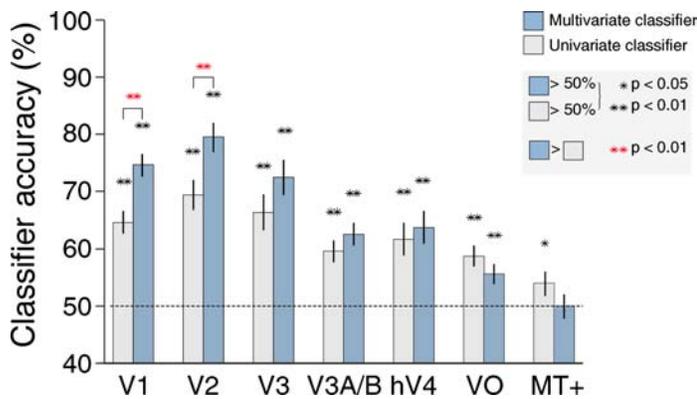
Figure 4. Univariate and Multivariate Classifier Performance across subjects. Mean univariate and multivariate classifier accuracy across subjects for each visual area are plotted in light gray and dark blue bars, respectively. Error bars indicate ±1 standard deviation of the population of classification accuracy estimates generated using the permutation analysis. Chance performance (50%) is shown with the dashed line. In all visual areas except MT+, the multivariate classification accuracy was significantly ($p < 0.01$, one-tailed permutation test) above the chance performance predicted by the permutation analysis, while the univariate classification accuracy was significantly above chance ($p < 0.05$, one-tailed permutation test) in all areas. The multivariate pattern classifier generally outperformed the univariate classifier, and this difference was significant ($p < 0.01$, two-tailed permutation test) for areas V1 and V2 when compared with the differences predicted by chance according to the permutation analysis. In areas VO and MT+, the univariate classifier outperformed the multivariate classifier, but this difference was not significant.

but classification performance on lime-magenta blocks were perfect, this would give an overall performance of 75%. We found that this was not the case; for both univariate and multivariate classifiers, classification performance for lime-magenta test stimuli and classification performance for orange-cyan stimuli showed a positive linear correlation (univariate: *slope* = 0.33, $R^2 = 0.13$, $p < 0.05$, multivariate: *slope* = 0.85, $R^2 = 0.61$, $p < 0.01$).

Increased performance of the multivariate classifier compared with the univariate classifier in V1 and V2 indicates that there are reliable, stimulus-related patterns of activity in these areas. If the pattern of activity across voxels were uninformative about the non-cardinal color of the stimulus, we would expect the multivariate classifier performance to be at best the same as the univariate case (since if the classifier learnt a pattern that was not stimulus-related, performance could decrease).

Discussion

We found evidence in human visual cortex for representations of color as early as V1 that combine information from the L–M and S–(L + M) opponent pathways

hypothesized to carry information in parallel from subcortical areas to cortex. The ability to use BOLD activity to discriminate stimuli matched for the postulated subcortical mechanisms demonstrates that the neural population must include neurons modulated by signals from both the chromatically opponent pathways. Below we discuss the implications of these results for how color is represented in human visual cortex, in particular the bias for lime-magenta over orange-cyan stimuli, and differences between visual areas in classifier performance.

Origin of asymmetry in the representation of two non-cardinal color modulations

We found a common bias across cortical visual areas for lime-magenta over orange-cyan blocks, even though our stimuli were matched for cone contrast, and for the response of sub-cortical pathways. The consistency of this bias across subjects suggests that it reflects a typical asymmetry in cortical representations of color. Specifically, this finding implies that there is a more numerous or more active population of neurons which respond to lime and/or magenta than to orange and/or cyan stimuli.

There is some evidence for a bias in the opposite direction in single-unit recordings in macaque V1, and from human psychophysics. Both Conway ([2001]) and Solomon and Lennie ([2005]) found a bias when testing the responses of macaque V1 cells to L, M and S-cone isolating stimuli. Of the 45 (Conway, [2001]) and 19 (Solomon & Lennie, [2005]) L–M color opponent cells that also responded to S-cone isolating stimuli, for 93% and 89% of cells (respectively) the response to the S-cone isolating stimulus had the same sign as the M-cone isolating stimulus; that is, the cells preferred a color direction closer to orange-cyan than lime-magenta. Krauskopf and Gegenfurtner ([1992]) report subtle psychophysical asymmetries for human observers between the non-cardinal axes in the effects of adaptation on discrimination threshold. Their data are consistent with a greater prevalence of adaptable mechanisms tuned to orange-cyan than to lime-magenta. Our data imply a bias of the opposite direction in human visual cortex, suggesting further work is necessary to reconcile these findings. In a recent study on human discrimination thresholds, Danilova and Mollon ([2010]) found that discrimination thresholds were lowest along a line in chromaticity space connecting unique blue and unique yellow. The hypothetical channel Danilova and Mollon ([2010]) propose to account for their results would lie closer to the lime-magenta modulation than the orange-cyan modulation, and these results could be a psychophysical correlate of the bias we observed in our fMRI data.

Organization of color processing in early human cortical areas

While significant classifier performance indicates representations of non-cardinal colors, differences in classifier

performance between areas are difficult to interpret. Brouwer and Heeger (2009) found highest classifier performance in V1, yet their principal components analysis suggested that the representation of color in hV4 and VO more closely matches our perceptual experience. Classifier performance depends not only on the presence of relevant information (here, non-cardinal representations of color) within an area, but also on the accessibility of this information at the coarse spatial scale of our functional measurements.

For areas V1 and V2, multivariate classifiers significantly outperformed univariate classifiers, showing that there was stimulus related information in the pattern of activity. In macaque V1 and V2 there are orderly maps of hue selectivity (including both cardinal and non-cardinal colors) across the surface of the cortex (Xiao, Casti, Xiao, & Kaplan, 2007; Xiao, Wang, & Felleman, 2003). If similar maps exist early in human visual cortex, their existence may increase the chance of biased sampling of chromatic preferences across voxels. The size of these hue maps, which each represent a large spectrum of hues, is only around 200 μm across the surface of the cortex in macaque V1, with individual maps separated by around 400 μm (Xiao et al., 2007). If maps of approximately the same size exist in human V1, a single voxel would contain approximately 6 hue maps. It is unlikely that a single voxel would sample neurons whose preference included only a narrow range of hues, but it could be that this map arrangement would make it more likely for biases between voxels to arise. Furthermore, in macaque, the hue maps in V2 are on average 2 to 2.5 times longer than the hue maps in V1 (Xiao et al., 2007). Larger maps with the same voxel resolution should increase the likelihood of biased sampling of hue maps, and increase the magnitude of the biases, which could underlie the tendency in our results for classifiers in V2 to outperform classifiers in V1.

Alternatively, it is possible that the stimulus related information in the pattern of activity is not due to qualitatively different patterns of response for lime-magenta and orange-cyan but instead a single pattern of visually responsive voxels which respond more strongly in the lime-magenta case. Since there is a univariate bias, the multivariate classifier could potentially be learning the difference between a strong signal in noise and a weaker version of the same signal (also in noise). The increased performance of the multivariate versus univariate classifiers might then be based purely on the ability of the multivariate classifiers to ascribe more to weight individual voxels on the basis of their signal-to-noise ratio. Further empirical and theoretical work will be required before it is possible to discriminate with certainty between these alternatives.

Response of dorsal visual areas

Poor classifier performance in MT+ is consistent with the classifier results of Brouwer and Heeger (2009), as well as evidence from MT of rhesus monkey (Britten, Shadlen, Newsome, & Movshon, 1992; Dubner & Zeki, 1971) and human MT+ (Huk, Dougherty, & Heeger, 2002; Tootell et al., 1995; Zeki et al., 1991) that this area is not generally selective for the color of surfaces and is less responsive to chromatic than achromatic stimuli (Gegenfurtner et al., 1994; Liu & Wandell, 2005; Wandell et al., 1999), although sensitivity to chromatic motion (Barberini, Cohen, Wandell, & Newsome, 2005; Wandell et al., 1999) has been reported. Likewise, the reduced classifier performance in V3A/B with respect to V1, V2 and V3 may reflect the general preference of this dorsal area for motion (Tootell et al., 1997), and its reduced responsivity to chromatically defined stimuli (Liu & Wandell, 2005). Nevertheless, for each subject, MT+ had fewer voxels than any other area we defined, which alone may account for decreased classifier performance.

Response of ventral visual areas

Areas hV4 and VO are often thought to be specialized for the processing of color. In macaque V4 evidence from both single unit recordings (Zeki, 1983) and from neuroimaging (Conway & Tsao, 2006) implicate this area as a 'color center'. In human, there is converging evidence from patients with cerebral achromatopsia (Zeki, 1990) and from PET (Lueck et al., 1989) and fMRI (Bartels & Zeki, 2000; Hadjikhani, Liu, Dale, Cavanagh, & Tootell, 1998; Liu & Wandell, 2005; McKeefry & Zeki, 1997; Mullen, Dumoulin, McMahon, de Zubicaray, & Hess, 2007; Wade, Brewer, Rieger, & Wandell, 2002) neuroimaging studies that both hV4 and VO are involved in color vision. Additionally, there is evidence that the response properties of VO match color perception in showing weaker responses to high than to low temporal frequencies (Jiang, Zhou, & He, 2007; Liu & Wandell, 2005), while V1 does not. It therefore might be expected that classifier performance would be greatest in these areas, but this was not the case for our stimuli, or for more perceptually relevant hues (Brouwer & Heeger, 2009). We consider five possible reasons for this.

First, our definitions of hV4 and VO may not include the region of ventral visual cortex that is specialized for color processing; Hadjikhani et al. (1998) reported color selectivity in area V8, but not hV4. We think this account of our results is unlikely because our definition of hV4, corresponding to that of Wandell et al. (2007), would include part of the V8 described by Hadjikhani et al. (1998), which showed color selectivity, with the remainder of their V8 corresponding to our VO.

Second, areas hV4 and VO may be more susceptible to task specific demands than other areas. We both asked subjects not to attend to the stimuli and required them to engage with a task at fixation. By diverting attention from the experimental stimulus we aimed to avoid artifactual classifier performance that was based not on differences in

the stimulus-driven response but on differences in attention between the two conditions. However, when attention is directed to a task unrelated to the stimulus, the stimulus-driven BOLD response is suppressed, and this suppression increases with the attentional load of the task (Rees, Frith, & Lavie, 1997; Schwartz et al., 2005). Single-unit recordings in macaque and fMRI in humans show that V4 and hV4 show greater attentional modulation than earlier visual areas (Hansen, Kay, & Gallant, 2007; Reynolds & Chelazzi, 2004; Schwartz et al., 2005).

Third, reduced performance of multivariate classifiers in hV4 and VO may result if our voxel size (1.5 mm each side) resulted in biases when sampling neural representations of color in V1 and V2, but not hV4 or VO. The spatial arrangement of chromatic preferences may be either less ordered, ordered in a different way, or ordered on a smaller spatial scale than in the earlier visual areas.

Fourth, any nonlinearity in the signal which differs between areas may enhance or reduce the stimulus discriminability in the BOLD response. For example, the contrast response of VO to L–M and S–(L + M) modulating stimuli is more nonlinear than V1 (Liu & Wandell, 2005). It is not clear how such nonlinearities should affect classifier performance, but it is possible that the reduced classifier performance in ventral areas were due to such nonlinearities.

Finally, there may be increased noise in our imaging results for these areas, which are typically located further from the surface of the head than the earlier visual areas and dorsal areas, and near the transverse sinus, which can cause imaging artifacts (Winawer, Horiguchi, Sayres, Amano, & Wandell, in press). We think it likely that the reduced performance of the classifiers in hV4 and VO reflects some combination of the impact of attention, nonlinearities, imaging noise and (for multivariate classifiers) the spatial arrangement of color processing within these areas, rather than implying their diminished selectivity for color.

Limitations of this study

All these conclusions are based on the assumption that our stimuli induced responses in sub-cortical pathways that are indistinguishable when considering the responses of each pathway independently. We consider a number of reasons why this assumption may be invalid.

Macular pigmentation selectively attenuates shorter wavelengths in the central two degrees of the visual field (Hammond, Wooten, & Snodderly, 1997; Wyszecki & Stiles, 1967). When defining our stimuli we used the Stockman and Sharpe (2000) 2 degree cone spectral sensitivities, which take into account the impact of macular pigmentation. Since macular pigmentation does not extend beyond the central 2 degrees of visual field (Hammond et al., 1997; Stringham, Hammond, Wooten, & Snodderly, 2006), it is possible that for the region

peripheral to this, our stimuli were no longer balanced for responses they induce in the sub-cortical pathways. To rule out this potential artifact, we repeated the classifier analysis, limiting the voxels included in the classifier to those within 2 degrees visual angle from fixation and thus excluding any voxels which respond to an area of visual field for which the stimuli may not have been balanced. With this analysis, we found classifier performance was reduced, but remained significantly above chance in all areas except MT+ (data not shown). This rules out the possibility that classifier performance in the original analysis was based on artifacts in the stimuli caused by macular pigmentation.

It is important also to emphasize the robustness of our conclusions to any inaccuracies in the determination of subjective equiluminance for each subject. For example, let us consider the situation if there was a 1% artifactual modulation in the lime-magenta blocks and no artifact in orange-cyan blocks. The 25% luminance modulation was added in opposite phases for different lime-magenta blocks, meaning that the effect of the luminance artifact would be to increase effective luminance contrast to 26% in half of the lime-magenta blocks and decrease it to 24% in the other half. For such a bidirectional effect to introduce a bias in the univariate response between lime-magenta and orange-cyan blocks, the contrast response function in the vicinity of 25% luminance contrast would have to be highly non-linear. Furthermore, any such bias would be unlikely to show the consistency between subjects observed here. In terms of the multivariate analysis, a classifier would have to learn a disjunctive discrimination between (24 or 26%) vs. (25%) modulation in luminance in order to be able to classify the stimuli based on luminance alone. Our use of linear classifiers minimizes the possibility that luminance artifacts were used as a cue to discrimination.

More fundamentally, it could be that the classic model of two chromatically opponent sub-cortical channels is insufficient to capture the selectivity of all neurons in the lateral geniculate nucleus (LGN). The majority of observed LGN responses can be accounted for by the classic model, but some evidence suggests that the chromatic responses of the LGN may not be fully described by the two channels. It is difficult to investigate physiological correlates of the higher-order color mechanisms in the LGN, since these mechanisms have primarily been revealed by psychophysical habituation (Krauskopf et al., 1986), and there is little or no habituation of cells in the LGN, yet these cells may contribute to the color tuning of the adaptable cortical cells (Tailby, Solomon, Dhruv, & Lennie, 2008a). Signals tuned to color directions away from the two opponent mechanisms have also been proposed (Webster & Mollon, 1991, 1994; Zaidi & Shapiro, 1993) that could theoretically be inputs to cortical areas. Finally Tailby, Solomon, and Lennie (2008b) recently reported that in macaque LGN, a subset of neurons responded to modulations along both the S–(L + M) and L–M axes. Our present study does not address the question of whether the model of two

chromatically opponent sub-cortical channels is sufficient to describe the coding of color by the LGN. We plan to investigate this possibility further in future investigations using human fMRI.

Appendix A

Multivariate classifier performance as a function of the number of voxels

For each subject, the classifier analysis was repeated for increasing numbers of voxels from each area. Figure A1 shows multivariate classifier performance as an increasing number of voxels were included in the classifier, on a

semilogarithmic scale. It also shows the best fitting curve (see Equation 1), from which the limit was reported as the classifier performance in Figure 4, or the mean classifier performance in cases where the curve did not fit the data.

Appendix B

Generalization of classification across different luminance–color pairings

Our main analyses, grouping lime-magenta and orange-cyan blocks of different luminance pairings, show that the classifier generalizes across luminance levels. This is
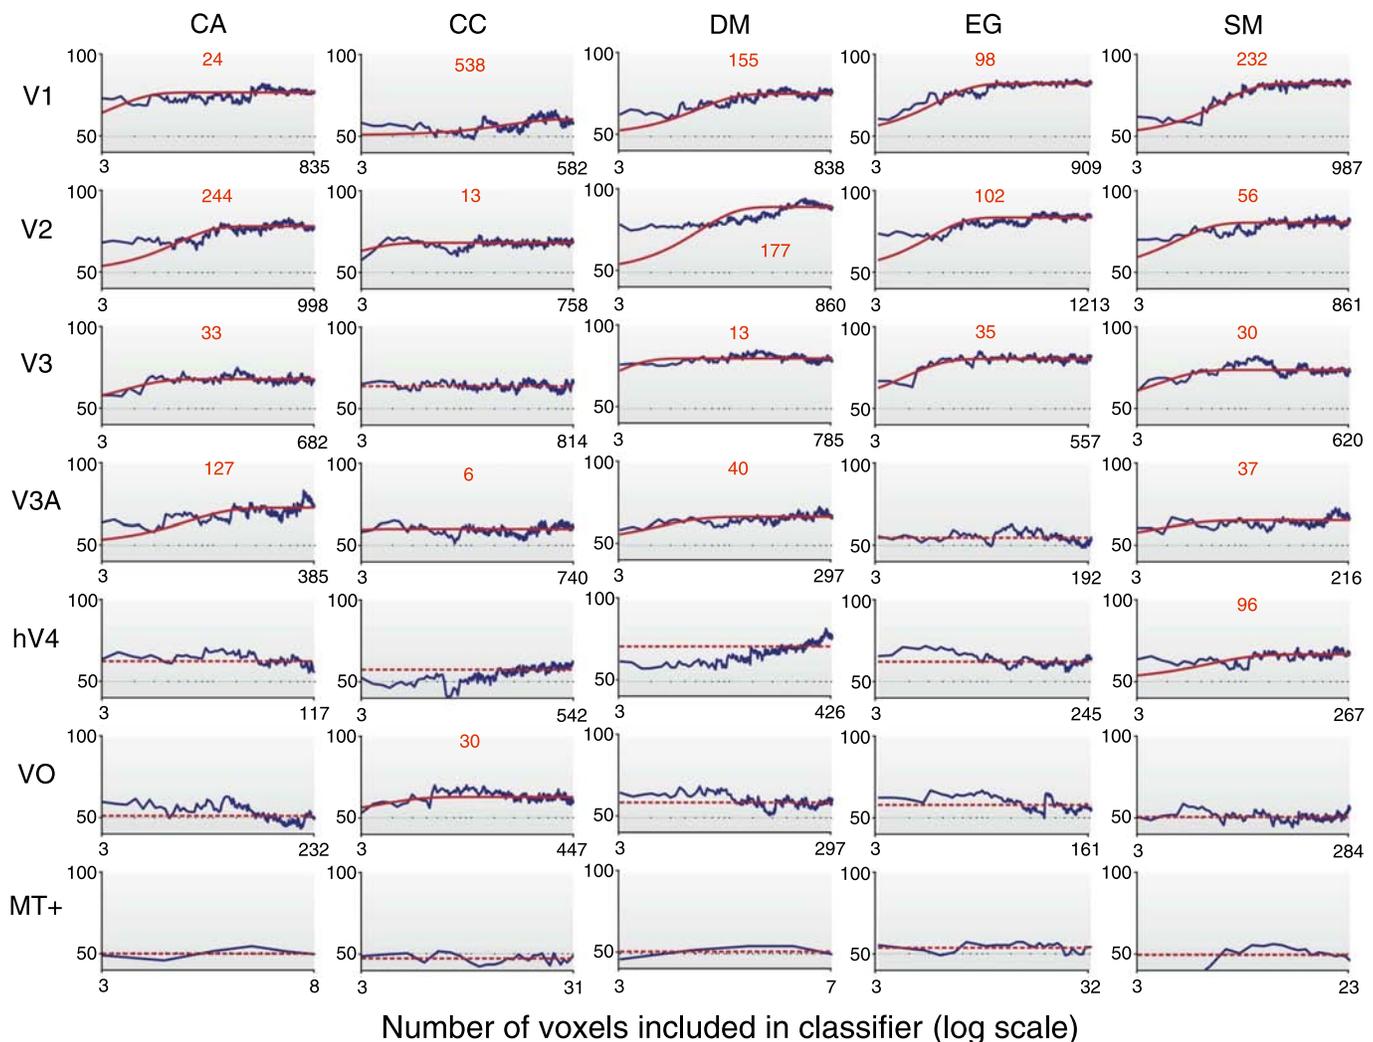

Figure A1. Classifier Performance as a function of the number of voxels included in the classifier, for each area and each subject. In each plot the filled blue line plots the classifier performance, and the filled red line plots the best-fitting exponential growth function, given by Equation 1. Where the exponential growth function fitted the data, the limit of this function was taken as the classifier performance in that area for that subject; the number of voxels taken to reach this limit is given by the red number on these plots. Where the exponential growth function did not fit the data (usually when performance was low), the mean was taken as the classifier performance, which is plotted as a dashed red line.
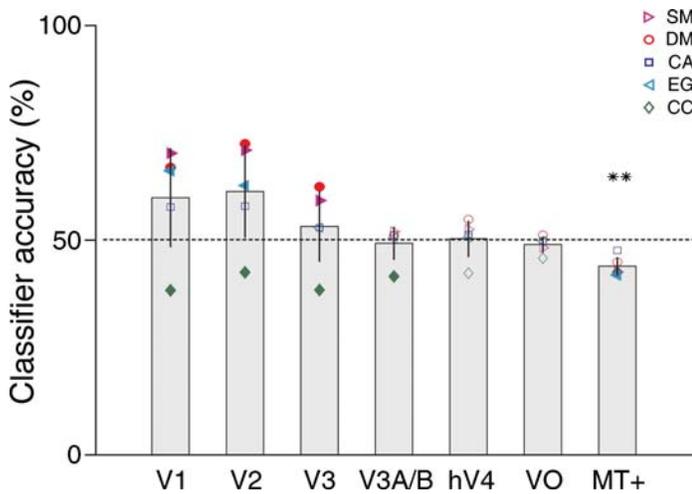
Figure B1. Multivariate classifier performance, training on one pair of luminance levels and testing on the other two. Results are averaged across the four different combinations of train/test pairs, error bars indicate the between subjects mean ±1 *SD*. The ** denote that the classifier was significantly different from chance ($p < 0.01$) in a two-tailed *t*-test. Additionally, the filled symbols indicate cases where the individual's classifier was significantly ($p < 0.05$) different from 50%, as assessed using a bootstrapped estimate of the variance by taking scoring the classifier's performance for 1000 iterations of randomly shuffled labels assigned to the test data.

because both training and test data contain an equal number of the two phases of luminance–color relations so the classifier has to generalize across different luminance levels to learn the stimuli. If the classifier can learn this more general rule, it at first sight seems reasonable to test whether it can generalize from one level to the other.

However, there is a subtle but important point to make here. If we train on one luminance level and test on the other, the classifier may learn a different rule than lime-magenta vs. orange-cyan. If the classifier is trained on dark lime-light magenta vs. dark orange-light cyan, the classifier could learn to separate the training data as dark green-light red vs. dark red-light green. If this different rule were learned by the classifier then when tested with light lime-dark magenta vs. light orange-dark cyan it would give the opposite result to learning the non-cardinal color modulation. Thus successful classification would provide evidence in support of the classifier learning the non-cardinal color modulation, and argue against the notion that the original classifier performance could have been based on an artifact. But a negative result neither supports nor excludes the possibility that classifier performance was due to a luminance artifact. In short, in the original analysis any small luminance artifact could have worked in favor of classifier performance but in this new analysis there is a 25% contrast luminance modulation working against classifier performance.

The results of this analysis are shown in Figure B1. We found that in area MT+, classification performance was significantly below chance ($p < 0.01$, two-tailed *t*-test) across subjects, and for individual CC in areas V1, V2, V3 and V3A/B ($p < 0.01$, permutation test). Below chance performance is consistent with the classifier being significantly good at learning a decision rule based on the pairing of luminance and one of the cardinal modulations (for example, successfully learning to separate the data as dark green-light red vs. dark red-light green). This below chance performance neither supports nor excludes the possibility that the original classifier performance was based on a luminance artifact.

For subjects SM, DM and EG classification performance was significantly ($p < 0.05$, permutation test) above chance for V1 and V2, with SM and DM also significantly ($p < 0.05$, permutation test) above chance in area V3. The fact that classifier performance is significantly above chance for some areas in some subjects gives evidence that in these cases the original classifier performance could not have been based on a luminance artifact.

Acknowledgments

References

Anstis, S., & Cavanagh, P. (1983). A minimum motion technique for judging equiluminance. In J. D. Mollon & R. T. Sharpe (Eds.), *Colour vision: Physiology and psychophysics* (pp. 156–166). London: Academic Press.

Barberini, C. L., Cohen, M. R., Wandell, B. A., & Newsome, W. T. (2005). Cone signal interactions in direction-selective neurons in the middle temporal visual area (mt). *Journal of Vision, 5*(7):1, 603–621, http://journalofvision.org/content/5/7/1, doi:10.1167/5.7.1. [PubMed] [Article]

Bartels, A., & Zeki, S. (2000). The architecture of the colour centre in the human visual brain: New results

and a review. *European Journal of Neuroscience, 12,* 172–193. [PubMed]

Bennett, K. P., & Campbell, C. (2000). Support vector machines: Hype or hallelujah? *SIGKDD Explorations, 2,* 1–13.

Birch, J., Barbur, J. L., & Harlow, A. J. (1992). New method based on random luminance masking for measuring isochromatic zones using high resolution colour displays. *Ophthalmic Physiology Optics, 12,* 133–136. [PubMed]

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436. [PubMed]

Britten, K., Shadlen, M., Newsome, W., & Movshon, J. (1992). The analysis of visual motion—A comparison of neuronal and psychophysical performance. *Journal of Neuroscience, 12,* 4745–4765. [PubMed] [Article]

Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *Journal of Neuroscience, 29,* 13992–14003. [PubMed]

Conway, B. R. (2001). Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). *Journal of Neuroscience, 21,* 2768–2783. [PubMed] [Article]

Conway, B. R., & Tsao, D. Y. (2006). Color architecture in alert macaque cortex revealed by fMRI. *Cerebral Cortex, 16,* 1604–1613. [PubMed]

Danilova, M., & Mollon, J. (2010). Parafoveal color discrimination: A chromaticity locus of enhanced discrimination. *Journal of Vision, 10*(1):4, 1–9, http://journalofvision.org/content/10/1/4, doi:10.1167/10.1.4. [PubMed] [Article]

Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *The Journal of Physiology, 357,* 241–265. [PubMed] [Article]

de Valois, R., Cottaris, N., Elfar, S., Mahon, L., & Wilson, J. (2000). Some transformations of color information from lateral geniculate nucleus to striate cortex. *Proceedings of the National Academy of Sciences of the United States of America, 97,* 4997–5002. [PubMed] [Article]

Dubner, R., & Zeki, S. M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research, 35,* 528–532. [PubMed]

Engel, S. A., & Furmanski, C. S. (2001). Selective adaptation to color contrast in human primary visual cortex. *Journal of Neuroscience, 21,* 3949–3954. [PubMed] [Article]

Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organization in human visual cortex and

the spatial precision of functional MRI. *Cerebral Cortex, 7,* 181–192. [PubMed]

Frackowiak, R. S., Friston, K. J., Frith, C. D., Dolan, R. J., & Mazziotta, J. C. (1997). *Human brain function.* San Diego: Academic Press.

Gegenfurtner, K. R., Kiper, D. C., Beusmans, J. M., Carandini, M., Zaidi, Q., & Movshon, J. A. (1994). Chromatic properties of neurons in macaque mt. *Visual Neuroscience, 11,* 455–466. [PubMed]

Hadjikhani, N., Liu, A. K., Dale, A. M., Cavanagh, P., & Tootell, R. B. (1998). Retinotopy and color sensitivity in human visual cortical area V8. *Nature Neuroscience, 1,* 235–241. [PubMed]

Hammond, B. R., Wooten, B. R., & Snodderly, D. M. (1997). Individual variations in the spatial profile of human macular pigment. *Journal of Optical Society of America A, Optics, Image Science, and Vision, 14,* 1187–1196. [PubMed]

Hansen, K. A., Kay, K. N., & Gallant, J. L. (2007). Topographic organization in and near human visual area V4. *Journal of Neuroscience, 27,* 11896–11911. [PubMed] [Article]

Haynes, J., & Rees, G. (2005a). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience, 8,* 686–691. [PubMed]

Haynes, J.-D., & Rees, G. (2005b). Predicting the stream of consciousness from activity in human visual cortex. *Current Biology, 15,* 1301–1307. [PubMed]

Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews on Neuroscience, 7,* 523–534. [PubMed]

Huk, A., Dougherty, R., & Heeger, D. (2002). Retinotopy and functional subdivision of human areas MT and MST. *Journal of Neuroscience, 22,* 7195–7205. [PubMed] [Article]

Ishihara, S. (1990). *Ishihara's tests for color-blindness, 38 plate ed.* Tokyo: Kanehara Shuppan Co. Ltd.

Jiang, Y., Zhou, K., & He, S. (2007). Human visual cortex responds to invisible chromatic flicker. *Nature Neuroscience, 10,* 657–662. [PubMed]

Joachims, T. (1999). Making large-scale SVM learning practical. In B. Scholkopf, C. Burges, & A. Smola (Eds.), *Advances in Kernel methods—Support vector learning* (chapter 11, pp. 41–56). Cambridge, MA: MIT Press.

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience, 8,* 679–685. [PubMed] [Article]

Kamitani, Y., & Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Current Biology, 16,* 1096–1102. [PubMed] [Article]

Kingdom, F., Moulden, B., & Collyer, S. (1992). A comparison between colour and luminance contrast in a spatial linking task. *Vision Research, 32,* 709–717. [PubMed]

Kourtzi, Z., Erb, M., Grodd, W., & Bülthoff, H. H. (2003). Representation of the perceived 3-D object shape in the human lateral occipital complex. *Cerebral Cortex, 13,* 911–920. [PubMed]

Krauskopf, J., & Gegenfurtner, K. (1992). Color discrimination and adaptation. *Vision Research, 32,* 2165–2175. [PubMed]

Krauskopf, J., Williams, D. R., Mandler, M. B., & Brown, A. M. (1986). Higher order color mechanisms. *Vision Research, 26,* 23–32. [PubMed]

Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F., & Baker, C. I. (2009). Circular analysis in systems neuroscience: The dangers of double dipping. *Nature Neuroscience, 12,* 535–540. [PubMed] [Article]

Larsson, J., & Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *Journal of Neuroscience, 26,* 13128–13142. [PubMed] [Article]

Liu, J., & Wandell, B. A. (2005). Specializations for chromatic and temporal signals in human visual cortex. *Journal of Neuroscience, 25,* 3459–3468. [PubMed] [Article]

Lueck, C. J., Zeki, S., Friston, K. J., Deiber, M. P., Cope, P., Cunningham, V. J., et al. (1989). The colour centre in the cerebral cortex of man. *Nature, 340,* 386–389. [PubMed]

Manjón, J. V., Lull, J. J., Carbonell-Caballero, J., García-Martí, G., Martí-Bonmatí, L., & Robles, M. (2007). A nonparametric MRI inhomogeneity correction method. *Medical Image Analysis, 11,* 336–345. [PubMed]

Mannion, D. J., McDonald, J. S., & Clifford, C. W. G. (2009). Discrimination of the local orientation structure of spiral glass patterns early in human visual cortex. *Neuroimage, 46,* 511–515. [PubMed]

McKeefry, D. J., & Zeki, S. (1997). The position and topography of the human colour centre as revealed by functional magnetic resonance imaging. *Brain, 120,* 2229–2242. [PubMed]

Mourao-Miranda, J., Bokde, A. L., Born, C., Hampel, H., & Stetter, M. (2005). Classifying brain states and determining the discriminating activation patterns: Support vector machine on functional MRI data. *Neuroimage, 28,* 980–995. [PubMed]

Mullen, K. T., Dumoulin, S. O., McMahon, K. L., de Zubicaray, G. I., & Hess, R. F. (2007). Selectivity of human retinotopic visual cortex to S-cone-opponent, L/M-cone-opponent and achromatic stimulation. *European Journal of Neuroscience, 25,* 491–502. [PubMed]

Parkes, L., Marsman, J., Oxley, D., Goulermas, J., & Wuerger, S. (2009). Multivoxel fMRI analysis of color tuning in human primary visual cortex. *Journal of Vision, 9*(1):1, 1–13, http://journalofvision.org/content/9/1/1, doi:10.1167/9.1.1. [PubMed] [Article]

Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. [PubMed]

Rees, G., Frith, C. D., & Lavie, N. (1997). Modulating irrelevant motion perception by varying attentional load in an unrelated task. *Science, 278,* 1616–1619. [PubMed]

Reynolds, J. H., & Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Reviews on Neuroscience, 27,* 611–647. [PubMed]

Schira, M. M., Tyler, C. W., Breakspear, M., & Spehar, B. (2009). The foveal confluence in human visual cortex. *Journal of Neuroscience, 29,* 9050–9058. [PubMed]

Schwartz, S., Vuilleumier, P., Hutton, C., Maravita, A., Dolan, R. J., & Driver, J. (2005). Attentional load and sensory competition in human vision: Modulation of fMRI responses by load at fixation during task-irrelevant stimulation in the peripheral visual field. *Cerebral Cortex, 15,* 770–786. [PubMed]

Seymour, K., Clifford, C. W. G., Logothetis, N. K., & Bartels, A. (2009). The coding of color, motion, and their conjunction in the human visual cortex. *Current Biology, 19,* 177–183. [PubMed]

Solomon, S. G., & Lennie, P. (2005). Chromatic gain controls in visual cortical neurons. *Journal of Neuroscience, 25,* 4779–4792. [PubMed] [Article]

Stockman, A., & Sharpe, L. T. (2000). The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research, 40,* 1711–1737. [PubMed]

Stringham, J. M., Hammond, B. R., Wooten, B. R., & Snodderly, D. M. (2006). Compensation for light loss resulting from filtering by macular pigment: Relation to the S-cone pathway. *Optometry Visual Science, 83,* 887–894. [PubMed]

Sumner, P., Anderson, E. J., Sylvester, R., Haynes, J. D., & Rees, G. (2008). Combined orientation and colour information in human V1 for both L–M and S-cone chromatic axes. *Neuroimage, 39,* 814–824. [PubMed]

Tailby, C., Solomon, S. G., Dhruv, N. T., & Lennie, P. (2008a). Habituation reveals fundamental chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience, 28,* 1131–1139. [PubMed]

Tailby, C., Solomon, S. G., & Lennie, P. (2008b). Functional asymmetries in visual pathways carrying S-cone signals in macaque. *Journal of Neuroscience, 28,* 4078–4087. [PubMed] [Article]

Tootell, R. B., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., et al. (1997). Functional analysis of V3A and related areas in human visual cortex. *Journal of Neuroscience, 17,* 7060–7078. [PubMed]

Tootell, R. B., Reppas, J. B., Kwong, K. K., Malach, R., Born, R. T., Brady, T. J., et al. (1995). Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *Journal of Neuroscience, 15,* 3215–3230. [PubMed]

Wade, A. R., Brewer, A. A., Rieger, J. W., & Wandell, B. A. (2002). Functional measurements of human ventral occipital cortex: Retinotopy and colour. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 357,* 963–973. [PubMed] [Article]

Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron, 56,* 366–383. [PubMed]

Wandell, B. A., Poirson, A. B., Newsome, W. T., Baseler, H. A., Boynton, G. M., Huk, A., et al. (1999). Color signals in human motion-selective cortex. *Neuron, 24,* 901–909. [PubMed]

Webster, M. A., & Mollon, J. D. (1991). Changes in colour appearance following post-receptoral adaptation. *Nature, 349,* 235–238. [PubMed]

Webster, M. A., & Mollon, J. D. (1994). The influence of contrast adaptation on color appearance. *Vision Research, 34,* 1993–2020. [PubMed]

Winawer, J., Horiguchi, H., Sayres, R., Amano, K., & Wandell, B. (in press). Mapping hV4 and ventral occipital cortex: The venous eclipse. *Journal of Vision.* [Article]

Wyszecki, G., & Stiles, W. S. (1967). *Color science: Concepts and methods, quantitative data and formulas.* New York: John Wiley & Sons, Inc.

Xiao, Y., Casti, A., Xiao, J., & Kaplan, E. (2007). Hue maps in primate striate cortex. *Neuroimage, 35,* 771–786. [PubMed] [Article]

Xiao, Y., Wang, Y., & Felleman, D. J. (2003). A spatially organized representation of colour in macaque cortical area V2. *Nature, 421,* 535–539. [PubMed]

Yushkevich, P. A., Piven, J., Hazlett, H. C., Smith, R. G., Ho, S., Gee, J. C., et al. (2006). User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage, 31,* 1116–1128. [PubMed]

Zaidi, Q., & Shapiro, A. G. (1993). Adaptive orthogonalization of opponent-color signals. *Biology Cybernics, 69,* 415–428. [PubMed]

Zeki, S. (1983). Colour coding in the cerebral cortex: The responses of wavelength-selective and colour-coded cells in monkey visual cortex to changes in wavelength composition. *Neuroscience, 9,* 767–781. [PubMed]

Zeki, S. (1990). A century of cerebral achromatopsia. *Brain, 113,* 1721–1777. [PubMed]

Zeki, S., Watson, J. D. G., Lueck, C. J., Friston, K. J., Kennard, C., & Frackowiak, R. S. J. (1991). A direct demonstration of functional specialization in human visual-cortex. *Journal of Neuroscience, 11,* 641–649. [PubMed]