Coding and Binding of Color and Form in Visual Cortex

Kiley Seymour[1,2], Colin W.G. Clifford[1,2], Nikos K. Logothetis[3,4] and Andreas Bartels[3,5]

[1]School of Psychology, University of Sydney, Sydney, NSW 2006, Australia, [2]Australian Centre for Excellence in Vision Science, Canberra, ACT 0200, Australia, [3]Max Planck Institute for Biological Cybernetics, 72076 Tübingen, Germany, [4]Imaging Science and Biomedical Engineering, University of Manchester, Manchester M13 9PT, UK and [5]Centre for Integrative Neuroscience, 72076 Tübingen, Germany

The processing of color and form is largely segregated within the visual brain. But there is also evidence to suggest that these features are coded in combination early in visual processing. Here, we combined high-resolution functional magnetic resonance imaging (fMRI) together with multivariate pattern classification to examine where in the visual cortex specific color form "conjunctions" are represented. Human subjects viewed visual displays containing colored spiral patterns. The spiral patterns could be red or green, and oriented either clockwise or counterclockwise, leading to 4 possible stimulus configurations. Two additional displays combined 2 of the above single color-form pairings, leading to double conjunctions. We applied linear classifiers to voxel activation patterns obtained while subjects viewed such displays. Our findings not only show that color and form information is coded across retinotopically defined visual areas, but also that the 2 double-conjunction stimuli can be distinguished. The voxels most informative about conjunctions were distinct from those most informative about color or form alone. Our results indicate that conjunctions of form and color may be coded by separate functional units as early as primary visual cortex. The results of this study have implications for theories concerning the segregation and binding of color and form information.

Keywords: binding, classifier analysis, color, fMRI, MVPA, orientation

Introduction

It is believed that early visual signals are separated into 2 chromatic channels with relatively poor spatial resolution underlying color perception (Derrington et al. 1984) and a third luminance channel with high spatial resolution that dominates form perception (Thorell et al. 1984; Liu and Wandell 2005). Electrophysiology and neurohistochemistry indicate that this early segregation of color and orientation information is represented cortically in the blobs and interblobs of primary visual cortex, respectively, (Livingstone and Hubel 1984; DeYoe and Van Essen 1985) and may persist throughout the visual cortex (Zeki 1978; Livingstone and Hubel 1988; Felleman and Van Essen 1991; Conway et al. 2007).

A visual system that processes color and form independently brings with it a binding problem: How can distributed processing result in a unified percept? Although there is evidence from macaque electrophysiology to suggest that some cells in early visual cortex may carry joint information about color and orientation (Thorell et al. 1984; Hubel and Livingstone 1990; Yoshioka and Dow 1996; Gegenfurtner et al. 1997; Kiper et al. 1997; Lennie 1998; Conway 2001; Johnson et al. 2001, 2008; Conway and Livingstone 2006), how these cells contribute to the perception of "bound" features is far from clear. For instance,

cells selective for orientation may exhibit residual chromatic tuning responses that have been inherited from the early segregation of cone inputs but are not necessarily "readout" in behavioral performance (Clifford et al. 2003; Forte and Clifford 2005; Peirce et al. 2008). Alternatively, some models of visual processing rely heavily on neural responses from multiplex cells for the perception of feature conjunctions of color and form (Billock 1995; Lennie 1998; Holcombe and Cavanagh 2001; Billock and Tsou 2004; Clifford 2010).

Most evidence of color-form conjunction coding in the human visual brain comes from studies using adaptation paradigms (McCollough 1965; Engel 2005). For instance, psychophysical investigations with the McCollough effect (McCollough 1965) have provided substantial evidence that some neurons may be tuned to specific combinations of color and orientation and play a role in the perception of bound features. If an observer adapts to, say, a red horizontal grating and a green vertical grating, then a subsequent achromatic test grating will appear in the complementary color: pale green if horizontal or pink if vertical. A lack of interocular transfer of the McCollough effect implies that it may be mediated at early stages of cortical processing (Barnes et al. 1999).

Adaptation paradigms in the context of functional magnetic resonance imaging (fMRI-A) have also been used to examine blood oxygen level–dependent (BOLD) responses to conjunctions of color and orientation (Engel 2005). As early as V1, Engel observed a component of adaptation specific to the combined color and orientation of the adapting stimulus over and above that which would be predicted from adaptation to the 2 attributes independently.

Although adaptation studies have been extensively used to provide evidence of selectivity to a given stimulus property, several reports have indicated that the tuning of adaptation may not be related in any simple way to the underlying neuronal selectivity for the stimulus property being examined (Abbott et al. 1997; Kohn and Movshon 2003; Hosoya et al. 2005; Tolias et al. 2005; Crowder et al. 2006). For instance, selective adaptation to stimulus properties, such as orientation and relative motion, have been observed as early as the ganglion cells of the retina even though response selectivity for these properties is not observed earlier than cortex (Hosoya et al. 2005). Hence, inferring the locus of neuronal response selectivity to feature conjunctions from the adaptation studies described above is limited by considerable uncertainty (Krekelberg et al. 2006; Sawamura et al. 2006; Clifford et al. 2007; Bartels et al. 2008).

In a recent study, Sapountzis et al. (2009) performed a direct comparison of multivariate pattern analysis (MVPA) and fMRI-A. They showed that MVPA was not only more sensitive in

detecting subtle orientation differences of stimuli but also that the voxels exhibiting high orientation preferences did not demonstrate significant adaptation to their preferred orientation.

Using the MVPA approach, Sumner et al. (2008) recently showed that decoding performance for orientations failed to generalize across L-M, S cone, and luminance channels. Although this could imply that conjunction coding of color and form arises early in the human visual cortex, observing a failure to generalize across these channels could simply reflect a residual chromatic tuning inherited from the early segregation of cone inputs and not necessarily a selectivity leading to the perceptual awareness of combined features. Furthermore, as generalization performance can only ever be expected to decline across a new data set, these results invite alternative explanations relating to differences in the quality of noise being presented to a classifier in the training and test sets. Specifically, a classifier may be forced to learn a pattern of responses from cue-invariant orientation detectors embedded in a background of unoriented, chromatically tuned responses. The optimal classifier weights for picking out the orientation signal would depend on the structure of the background response, which would in turn depend on the chromatic signature of the unoriented mechanisms stimulated. Thus, in the study of Sumner et al. (2008), one might expect a failure to generalize orientation discrimination performance across stimulus color even if the responses of the orientation detectors being sampled were actually color invariant.

In the current study, we used what we consider to be a more direct approach to investigating conjunction coding of color and form. We examined whether BOLD signals in human visual cortex could correctly discriminate stimuli that differed only by their specific pairing of color and orientation, using double-conjunction stimuli. This approach avoided using supralinear adaptation or decreases in generalization performance to infer conjunction specificity. Our results show that conjunctions of color and orientation can be decoded from patterns of voxel activation as early as V1. In addition they show that the most informative "conjunction" voxels do not overlap with voxels most informative about color or orientation alone. This suggests that separate functional units may be tuned to specific combinations of color and orientation with a likely role in the perceptual awareness of bound features.

Materials and Methods

Subjects

Five subjects (one author, K.S.) participated in this study. All had normal, or corrected to normal, visual acuity, and color vision. Each subject was familiarized with the stimuli and the task outside of the scanner and gave informed consent according to the joint ethics committee of the Max Planck Institute and University of Tübingen.

Stimulus

Avoiding Radial Biases from Oriented Stimuli

Typically, the stimuli used to address orientation selectivity are oriented gratings. It has been shown recently, however, that gratings of different orientation generate different levels of activation across retinotopic cortex (Furmanski and Engel 2000; Sasaki et al. 2006). Firstly, gratings oriented cardinally (vertically and horizontally) have been reported to generate larger responses in primary visual cortex than oblique gratings (Furmanski and Engel 2000). Secondly, although oblique gratings tilted clockwise and counterclockwise from vertical

generate comparable levels of mean activation across the visual field in early cortical areas, the response in any given quadrant representation of the visual field is consistently higher for gratings oriented close to radial than for the perpendicular orientation (Sasaki et al. 2006). The use of conventional grating stimuli may therefore confound studies using multivariate approaches to investigate orientation processing. To minimize the possibility of such biases in our study, we used as stimuli a pair of orthogonal "spirals" (Mannion et al. 2009). Each spiral had the property that everywhere in the stimulus the local orientation was ±45 deg to radial. The rationale for the use of these spiral gratings was that the amount of local orientation information was comparable with that conveyed by conventional gratings but that both cardinal and radial response biases in retinotopic visual areas were avoided (i.e., the clockwise- and counterclockwise-oriented spiral stimuli were matched in all regards at any visual location apart from their handedness).

Conjunction and Double-Conjunction Stimuli

In total, we presented 6 stimulus conditions in a blocked design, all consisting of colored spiral gratings presented in an annulus extending from 2 to 9 deg of eccentricity (i.e., width of ring subtended 7 deg). Four single-conjunction stimulus conditions resulted from the 4 possible combinations of pairing 1 of 2 colors with 1 of 2 spiral gratings: red clockwise (RCW), red counterclockwise (RCCW), green clockwise (GCW), and green counterclockwise (GCCW). We also created 2 double-conjunction conditions (RCW&GCCW or RCCW&GCW), each comprising 2 single color-orientation pairings that alternated at a frequency of 1 Hz within a block (the order of which was counterbalanced across runs). The significance of using double conjunctions is that the 2 conditions both contain the same 2 colors and orientations but differ exclusively in the conjunction of those features (see Fig. 1). As linear classifiers integrate evidence provided (separately) by each voxel about category membership, the logic here is that a classifier will only show above-chance decoding performance if some voxels are individually sensitive to the dimension of interest, that is, conjunctions of color and orientation. Hence, the assumption is that a classifier cannot rely on segregated color and orientation signals to distinguish the 2 double-conjunction conditions because in such a case both conditions would elicit the same summed neural responses within a block and thus result in chance-decoding performance.

It is important to note here that this logic has been previously used in a similar study to obtain explicit evidence for conjunction coding of color and motion under conditions of motion transparency (Seymour et al. 2009). In the current study, however, the use of double-conjunction stimuli comprising temporally interleaved single conjunctions (as opposed to a transparent superposition) affords a possible alternative interpretation of above-chance classification performance. The distinction is an important one, as each double-conjunction stimulus not only differed in the conjunctions present within a block but also potentially by the temporal profile of neural responses resulting from the alternation of single conjunctions. Therefore, although linear summation of separate color and orientation signals would elicit the same responses over a block, we cannot exclude the possibility that a classifier could distinguish double-conjunction stimuli based on a mixing of color-specific with orientation-specific neural responses together with a neurovascular nonlinearity (e.g., BOLD saturation within a voxel). Although the transparent superposition of single-conjunction stimuli might have circumvented this limitation, piloting of such stimuli resulted in qualitatively different arrangements of green–yellow and red–yellow boundaries within each double-conjunction stimulus that tended to result in monocular rivalry and would have made classifier performance difficult to interpret. See also the discussion on this point.

Basic Stimulus Parameters

Each stimulus consisted of square-wave colored spiral gratings (20 cycles) on a black background (mean stimulus luminance: 5 cd/m^2). The local orientation everywhere within each grating was ±45 deg to radial. Colors were set to isoluminance with a gray patch of 10 cd/m^2 using the minimum flicker technique inside the scanner for each subject (Kaiser 1991). Hues were adjusted such that they combined to gray, thus ensuring matched saturation and color-channel load. In
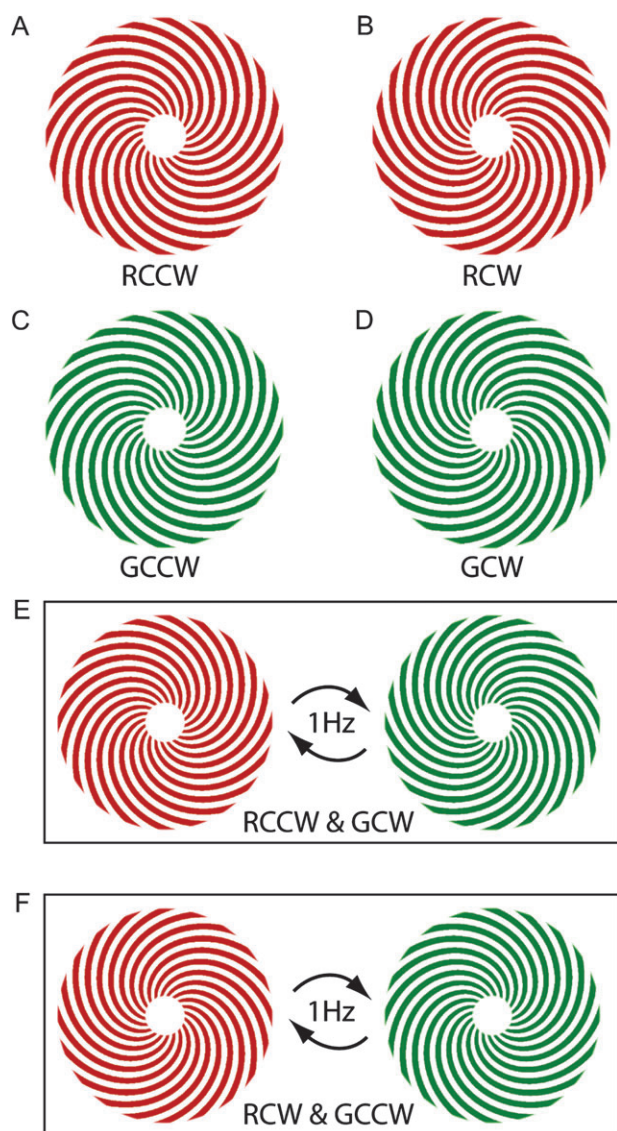
Figure 1. Schematic representation of the 6 stimulus conditions used in the experiment. (*A–D*) Single-conjunction stimuli, containing 1 color and 1 spiral orientation pattern each. Conditions *E* and *F* are double-conjunction stimuli, each containing both colors and both orientations, alternating between 2 single conjunctions at any one time. Note that the double-conjunction conditions would be indistinguishable to a linear classifier without conjunction-specific responses. Thus, a classifier could not rely on activation from color-specific and form-specific voxels to distinguish the stimuli. As all 4 features are displayed in both conditions.

addition to this, we added luminance noise by presenting colors at 1 of 2 luminance settings (9 or 11 cd/m^2) randomly in each block. Each condition was presented under each of these luminance settings the same number of times within a run. Blocks of the 4 single-conjunction stimuli were presented randomly as either "high" or "low" in luminance, and double-conjunction blocks could be low-low, low-high, high-low, or high-high for the 2 colors. This approach minimized the possibility of artefactual luminance-based color classification. To ensure that classification could not be based on the spatial representation of stimuli in retinotopic regions, each spiral grating was displayed in 1 of 2 phases. For example, on every second presentation, the previous colored stripes were black and vice versa. All stimuli were presented using Cogent2000 version 1.27 (Romaya, Wellcome Department of Imaging Neuroscience, London; http://www.vislab.ucl.ac.uk/cogent_ graphics.php) running under Matlab 2006b (Mathworks, Inc.). They were presented at a resolution of 1280 × 1024 pixels and at a screen refresh rate of 75 Hz at a viewing distance of 82 cm on a screen

subtending 24 × 18 visual degrees. Stimuli were projected onto a screen at the end of the scanner bore and viewed through a tilted mirror fixed to the head coil.

Task
Blocks of each condition lasted 12 s and were presented 5 times throughout 1 run of the experiment in pseudorandom permutations of the 6 conditions. There were a total of 8 runs, each lasting 6 min. Each block type was preceded equally often by the others.

In order to ensure that subjects were attending to the stimuli, we incorporated changes in luminance (30% decrements or increments) at 2 random intervals throughout each block, each lasting 500 ms. As subjects perceived a change, they indicated via a button press whether the change was a luminance increment or decrement. These manipulations were balanced across stimulus conditions. Subjects were required to fixate on a central fixation cross while carrying out this task.

fMRI
Functional images were acquired in a 3-T Siemens (Erlangen, Germany) TIM scanner using a gradient echo planar imaging sequence and a 12-channel phased-array head coil. We collected 27 slices positioned over the visual cortex, using an interleaved sequence with the following parameters: repetition time 2.34 s; echo time 39 ms, 96 × 96 matrix; and voxel size 2 × 2 × 2 mm. A high-resolution (1 mm isotropic) T1-weighted 3D-MDEFT image was acquired for surface reconstruction and used as an anatomical reference.

Preprocessing
We carried out the minimum of preprocessing using BrainVoyagerQX. Data were coregistered in raw AC–PC space and not transformed to any standard coordinate system. We corrected for head motion and made a mean intensity adjustment (global scaling); no smoothing was applied.

Retinotopic Mapping and Localizer Stimuli
We used phase-encoded retinotopic mapping, according to standard procedures (Sereno et al. 1994; Wandell et al. 2000). To each subject, we presented 3 rotating wedge runs and 2 expanding ring runs lasting 6 min each. For the delineation of retinotopically mapped visual areas using the phase-encoding method, cortical inflations of each subject were reconstructed from a high-resolution T1-weighted image. Gray and white matter was segmented and the cortex was reconstructed using BrainVoyagerQX (Kriegeskorte and Goebel 2001). A linear correlation between wedge position and neural activity was performed for each voxel, and borders of areas V1–hV4 and V3A/B were delineated manually on the basis of field sign alternations. V3A/B was defined as the retinotopically organized region dorsal to V3d representing both lower and upper visual field, we did not separate V3A from V3B. The hV4 complex was defined as voxels anterior to V3v in the fusiform gyrus representing both upper and lower visual fields (Brewer et al. 2005). In addition, we localized regions generally responsive to our stimuli using a 6-min run of alternating blank fixation and double-conjunction blocks of 12 s each. Those voxels that reached a liberal threshold of at least $P < 0.01$ (uncorrected) and fell within the above retinotopically defined regions of interest (ROIs) were selected to be used in the multivariate analysis.

Multivariate Analysis
We used an MVPA approach similar to that of earlier studies (Haynes and Rees 2005; Kamitani and Tong 2005). Multivariate analyses were carried out separately for each subject for each ROI. We included all voxels of a given ROI that reached threshold activation in response to the localizer stimuli. As decoding performance is heavily dependent on the spatial layout of feature information in a ROI (see, e.g., Bartels et al. 2008), and as differences in the size of each ROI might affect overall performance of a classifier, we refrained from comparing classification performance across ROIs.

For each ROI, the data were analyzed using the general linear model which fitted a separate regressor to every stimulus block. Each regressor was a 12-s boxcar convolved with the canonical hemodynamic response function. The amplitudes of the weights assigned to each regressor (beta estimates) by the model therefore conveyed how strongly a voxel

responded to each stimulus block (Friston et al. 1995). For every voxel this resulted in 40 beta estimates (8 runs × 5 presentations) for each of the 6 conditions. The beta estimates were then used to train linear support vector machines (SVMs) to distinguish responses to orientation, color, and double conjunctions (Burges 1998; Vapnik 1998). We used the Matlab implementation of SVMs provided by the Spider toolbox to achieve this (http://www.kyb.tuebingen.mpg.de/bs/people/spider/).

Prior to training SVMs, beta estimates from each voxel were normalized to a mean of 0 and a standard deviation (SD) of 1 across all trials of the 2 conditions. Outliers were removed from the data by setting all values that were beyond 2SDs from the mean to a fixed value of ±2SD. A leave-one-out approach was employed, in which a classifier was trained on voxel responses from $N - 1$ blocks and tested on the remaining block, cycling through all blocks. For each ROI of each subject, the mean classification accuracy of the test blocks was calculated (chance performance was always 50%). Significance was determined for each ROI using a one-sample t-test performed across the classification accuracies of the 5 subjects.

Examining Color and Orientation Selectivity
We pooled across single-conjunction conditions to train SVMs for color or orientation discrimination. For discrimination of color, we pooled blocks of RCW with RCCW and GCW with GCCW; for discrimination of orientation, we pooled RCW with GCW and RCCW with GCCW. Thus, when training an SVM for 1 feature (e.g., color discrimination), the SVM had to generalize over the other (e.g., orientation).

In addition to this analysis, our stimuli also allowed us to examine the extent to which each visual area would bias classification toward orientation or color information. To achieve this, we first trained an SVM to discriminate, for example, RCW versus GCCW and then tested on an entirely different data set consisting of GCW versus RCCW. Classification would therefore result in 1 feature (e.g., color) being correctly predicted, whereas the other (e.g., orientation) would be predicted incorrectly. The algorithm could thus either classify color correctly and get orientation wrong or vice versa. The result would reveal which of the 2 features was more "salient" to the classifier, for each region. Note that as there was no a priori baseline for this bias analysis, the absolute classification performance for each region was less informative than the differences between regions, as, for example, for this particular stimulus orientation may have affected BOLD responses in all ROIs more strongly than color. Therefore, differences in the strength of this bias would nevertheless be informative about distinct properties of the ROIs.

We also examined whether orientation could influence color classification performance and, conversely, whether color had an influence on orientation classification. In other words, if a classifier was trained to discriminate color, but only on the basis of CW oriented trials, would it perform equally well when tested on color discrimination across CCW trials? To test this, we trained an SVM to discriminate RCW versus GCW (39 blocks) and tested on the remaining block. We then also tested iteratively on each block from a data set consisting of RCCW versus GCCW. Equal performance in both cases would suggest that color classification was unaffected by orientation. The same procedure was used to examine orientation classification and its dependence on color.

Testing for Conjunction Selectivity
The primary aim of this study was to recover evidence for the presence of information encoding feature conjunctions across the examined visual areas. We used the 2 double-conjunction conditions (RCW&GCCW vs. RCCW&GCW) to train and test an SVM. Blocks from each category contained both colors and both orientations, and therefore, successful classification relied on distinct voxel pattern responses evoked solely by the 2 distinct "conjunctions" within these stimuli.

Voxelwise Segregation of Color, Orientation, and Conjunction Information
In order to test whether information about color, orientation, and conjunctions was conveyed by the same or different sets of voxels, we compared the absolute (i.e., unsigned) magnitudes of the classifier weights that had been assigned to each voxel during training on the distinct stimulus categories. Thus, a given (absolute) weight map of n voxels can be considered a vector in n-dimensional space. The angle

between 2 such weight vectors obtained from separately trained SVMs indicates their similarity. For each pair of weight maps, we calculated their angle. In addition, we also calculated the mean angle resulting from 1000 random permutations of their values that yielded the expected angle, given the assumption of no relation between the 2 vectors. For each ROI, we then tested (t-test) whether across the 5 subjects the actual angles between 2 weight maps differed significantly from the expected angle. We carried out the analysis using only voxels whose weights exceeded ± 2SD for either of the features, to test whether the highly informative voxels of 2 differently trained SVMs were more (or less) related than expected by chance. Significantly smaller angular differences than expected by chance would indicate a positive relation, in that voxels informing on, for example, orientation, were also informative about color, etc. Significantly larger angular differences than chance would indicate a negative relation, in that the better a given voxel coded for example, orientation, the worse it coded, for example, for color. A lack of significant angular difference would imply that weights, for example, for color were not related to those, for example, orientation. To ensure that weight maps reflected the voxel sensitivity purely for a given feature (e.g., color) in this analysis, we trained classifiers to discriminate colors only on CW trials (or only on CCW trials), such that the classifier would not be forced to generalize across orientation. Likewise, orientation weights were obtained from training on uniformly colored trials.

Results

Univariate Analysis
A conventional whole-brain univariate analysis of each subject revealed no systematic bias in activation to a particular spiral orientation, color, or conjunction condition.

Of particular interest were biases to spiral orientations, as we used spiral stimuli specifically to avoid the large-scale retinotopic biases reported for oriented grating stimuli (Sasaki et al. 2006). Although inevitably we found some voxels that showed preferences for one spiral orientation versus another at a relaxed uncorrected threshold of 0.05, the distribution of activity did not show a consistent retinotopic organization across subjects. Furthermore, no voxel survived a Bonferroni correction ($P < 0.05$) over the whole brain. Specifically, our spiral stimuli did not produce biases specific to particular quadrant representations of early visual areas, such as those reported by Sasaki et al. (2006) for oriented gratings.

As the primary aim of this study was to decode color-orientation conjunctions, it is important to note that even if there were consistent biases evoked by spiral orientations, these could not affect decoding results for feature conjunctions: Each of our 2 double-conjunction stimuli contained both orientations and would thus elicit the same bias. Hence, decoding must rely solely on the specific color orientation conjunction responses. Furthermore, as no net bias for a particular double-conjunction stimulus was observed in any ROI for any subject (Supplementary Fig. 1), and no reliable classification performance could be obtained by training on the mean response amplitude of each ROI (data not shown), large-scale spatial inhomogeneities at the univariate level could not explain subsequent conjunction classification performance at the multivariate level.

Decoding Orientation and Color
We first tested whether SVMs could discriminate between voxel activation patterns evoked by stimuli containing 2 different spiral orientation patterns. Note that we pooled trials across color, such that the classifier learned to distinguish

orientations regardless of their color (see Materials and Methods). Across all subjects, orientation was decoded in all areas examined, with mean performance ranging from 71.2% to 90.3% across ROIs. Each area performed better than the chance level of 50% at a significance of $P < 0.005$ in a t-test applied across 5 subjects (Fig. 2a). V1 and V2 were highest performers (89.1% and 90.3%).

Secondly, we trained SVMs to distinguish between the 2 colors. Here, we pooled trials across orientations, such that classifiers learned to distinguish colors regardless of their orientation. All areas of visual cortex achieved significant color discrimination across all subjects (Fig. 2b). Mean prediction accuracy ranged from 64.2% to 74.3%. hV4 was the highest performer (74.3%).

A Bias toward Color over Orientation in hV4

Apart from decoding color and orientation from each ROI, we also wanted to learn which of the 2 features was represented more strongly in each of the ROIs. We therefore ran an analysis to address whether each visual area would rely more on orientation or color information (see Materials and Methods). Figure 3 reports decoding biases for each ROI in terms of percent correct orientation classification (50% indicating no bias, greater than 50% signifying a bias toward orientation, and less than 50% signifying a bias toward color). With the exception of hV4, all areas in the visual cortex favored a correct classification of orientation rather than color (ranging from 62.4% to 66.2%, $P < 0.05$, $n = 5$). Area hV4 was the only region observed with a trend (46.0%, $P = 0.11$ [2-tailed t-test, $n = 5$]),
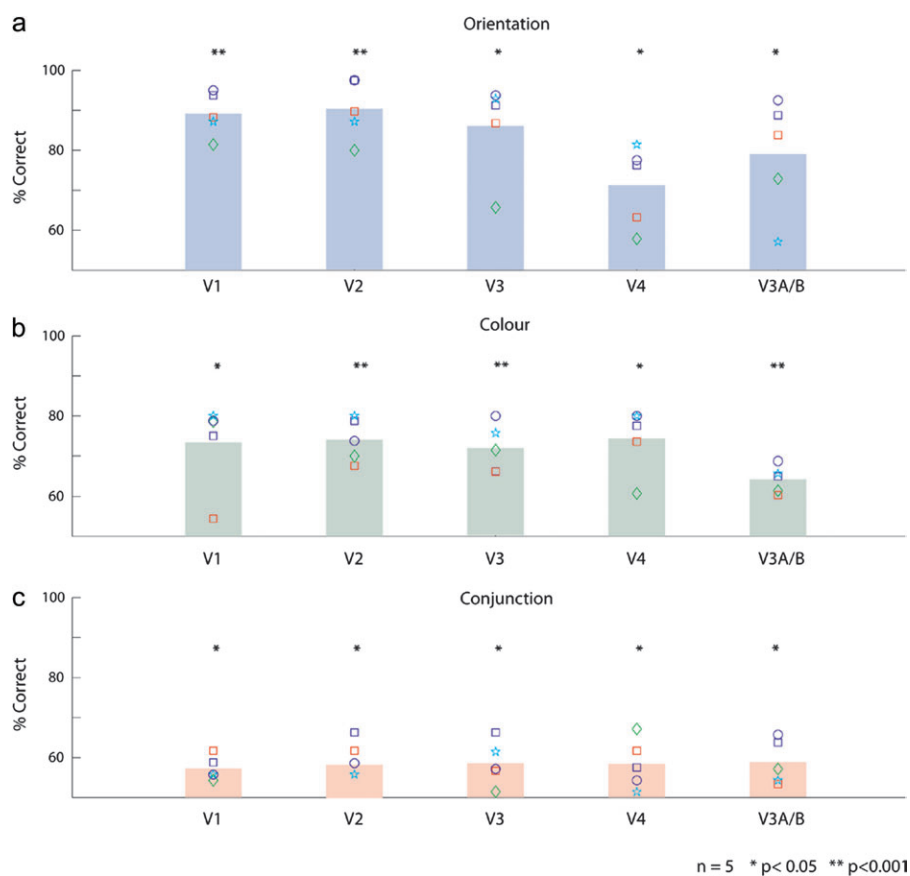
Figure 2. Discrimination performance of a linear classifier (SVM) at orientation, color, and double-conjunction stimuli. (a) Decoding of orientation (CW vs. CCW) was significantly above chance (chance = 50%) in every ROI, averaged across 5 subjects. (b) Decoding of color (red vs. green) was also significant in each area (c) decoding of double-conjunction stimuli (GCW&RCCW vs RCW&GCCW) reached significance in all areas. Symbols: single subject–decoding performances (n = 5 subjects).
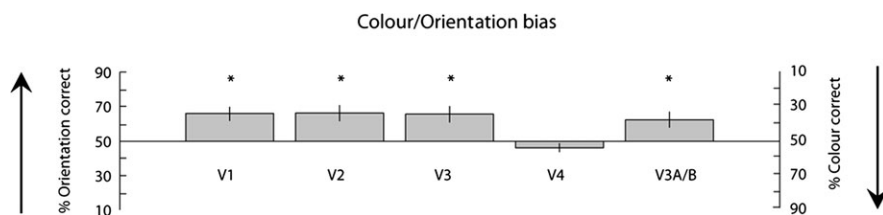

Figure 3. Classification bias to color or orientation in visual cortical regions. The classifier was trained on voxel responses to one set of single-conjunction stimuli, for example, RCW versus GCCW and tested on another, for example, GCW versus RCCW. The classifier was thus forced to achieve a correct orientation classification performance at the expense of an incorrect color classification or vice versa (50% = chance, i.e., the classifier relied equally often on color and orientation information). All areas except for hV4 showed a significant bias toward orientation classification. Classification based on hV4 (biased toward color) differed significantly from all other areas (biased toward orientation) ($P < 0.05$). Error bars: standard error (SE) across 5 subjects.

albeit nonsignificant, to favor the classification of color over orientation. Note that the above absolute biases are not necessarily meaningful, because, for example, a bias toward orientation decoding may not necessarily imply stronger selectivity for orientation in a region, as it could equally reflect the relative signal strength in our "stimuli" and possible biases in attention to orientation versus color. Therefore, in this analysis, the more meaningful analysis is to compare biases between ROIs, as only this would reveal ROI-specific feature biases independent of stimulus- or attention-driven overall biases. We therefore compared feature-decoding biases across areas in a one-way analysis of variance. Here a significant difference emerged in favor of color between area hV4 and all other ROIs ($P = 0.01$, $n = 5$).

Invariance of Color to Orientation Decoding and Vice Versa

Finally, we wanted to test to what extent decoding of one feature was invariant with respect to the other feature (see Materials and Methods). Specifically, we examined to what extent color coding was invariant with regard to orientation (Fig. 4a) and to what extent orientation coding was invariant with respect to color (Fig. 4b). In each analysis, decoding performance remained significantly above chance. Furthermore, the results showed that color did not have a significant influence on orientation decoding performance in any visual area (Fig. 4a). In other words, when a classifier was trained to discriminate 2 orientations but only on green stimuli, it performed equally well in decoding orientations from red stimuli as from green stimuli. Note that this ability to generalize across color—which in this study did not segregate any particular color channel—suggests that previous findings of conjunction coding in V1 based on a failure to generalize (Sumner et al. 2008) may simply reflect a residual chromatic tuning of orientation responses or the differences in the quality of noise that was presented to a classifier during the training and test phase. In addition, we found no significant difference in color-decoding performance when an SVM was forced to generalize color predictions across orientations (Fig. 4b).
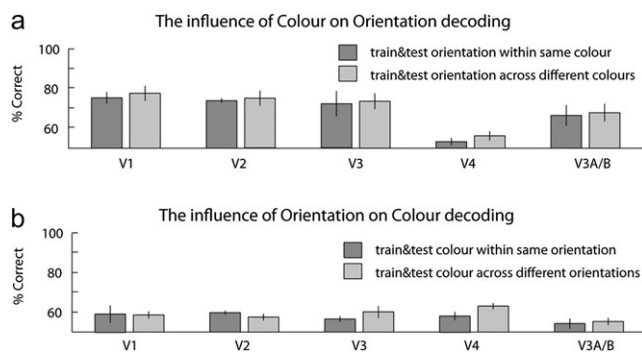

Figure 4. Invariance of orientation coding with respect to color and vice versa. (a) The classifier was trained to discriminate responses to orientation stimuli presented in one color, and then tested on responses to orientation presented in the same or different color. Classifier performance was not affected when training and test stimuli were different colors. (b) Color discrimination: the classifier was trained to discriminate responses to red or green stimuli in one orientation and then tested on responses to color stimuli presented in the same or opposite orientation. Classifier performance was not affected when training and test stimuli had different orientations. Error bars: SE across 5 subjects.

Decoding Conjunctions

We directly assessed conjunction coding of color and orientation using stimulus blocks that each contained 2 conjunctions (RCW&GCCW vs. RCCW&GCW trials). Each stimulus block contained both clockwise and counterclockwise orientations and both red and green color, the critical difference being the stimulus-specific conjunctions of these features (see Fig. 1). SVMs therefore had to distinguish one double conjunction from another, with all basic features being exactly matched in the 2 conditions. We found that SVMs could discriminate responses to distinct double conjunctions in all areas of visual cortex examined. Mean prediction performance ranged from 57.2% to 58.8%, reaching significance across all subjects in every ROI ($P < 0.05$, $n = 5$) (Fig. 2c).

Feature Segregation at the Scale of Voxels

Given that all visual areas contained information about color, orientation, and even about specific orientation–color conjunctions, we were interested to learn whether it was the same voxels or different sets of voxels that contained information about these distinct features. We tested this by examining to what extent the subset of voxels that obtained high weights in SVMs trained to discriminate, for example, orientation had also obtained high weights in SVMs trained to discriminate, for example, color. Our quantification compared the angles between weight vectors obtained from the 3 classifier types and tested whether these differed significantly from those expected by chance, for each ROI of each subject (see Materials and Methods). In all ROIs, this analysis found that the more reliable a voxel was at coding color, the worse it was at discriminating orientation (in radians: mean angle difference to that expected by chance: 0.07 ± 0.02SD, max = 0.09, min = 0.01, $n = 25$ (5 subjects × 5 ROIs), $P < 0.05$ in every ROI). Furthermore we found similar negative relationships ($n = 5$, $P < 0.05$) between conjunction-informative voxels and those most informative about orientation in V1 (mean: 0.09 ± 0.04SD), V2 (0.07 ± 0.06), V3 (0.07 ± 0.06), and V4 (0.07 ± 0.04). No significant relationship between these 2 features was observed in V3A/B. When weight vectors for conjunctions were compared with those of color, a negative relationship ($n = 5$, $P < 0.05$) was observed in a number of ROIs including V1 (0.06 ± 0.04), V2 (0.09 ± 0.05), and V3 (0.1 ± 0.05), indicating that the better a voxel coded for conjunctions the worse it coded for color. No significant relationship was observed in V3A/B or hV4. These findings indicate that, in general, the strength with which voxels code for one feature is anticorrelated with the strength with which they code for other features. Importantly, this also applied for conjunctions. Despite this, the best-coding voxels for one feature may (at least in some ROIs, data not shown) still contain information that is sufficient to decode other features. Nonetheless, these results indicate the voxels most informative about conjunctions are distinct from the voxels most informative about color or form in early visual cortex. Figure 5 shows a map of V1 illustrating the spatial distribution of the most informative voxels coding for color, orientation, and conjunctions in subject KS.

Discussion

We used MVPA to investigate whether information about color, orientation, and specific conjunctions of these 2 features could be decoded from patterns of BOLD response recorded in the
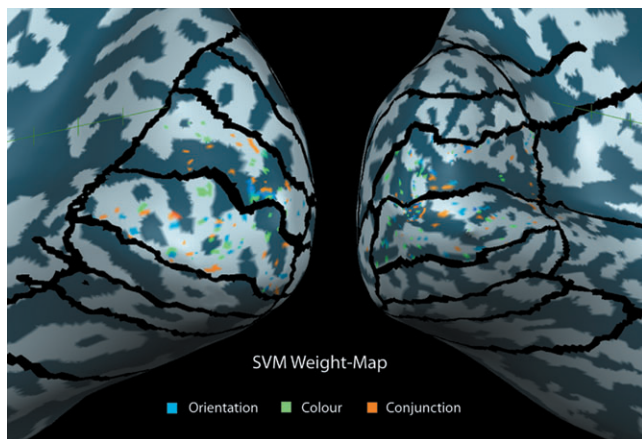
Figure 5. A weight map of the best-coding voxels (≥2SD) in V1 for subject KS as determined by 3 SVM analyses, for orientation, color, and conjunction discrimination. Note their apparent random (nonoverlapping) distribution for each condition and the lack of hemifield or quadrant biases.

human visual brain. We found evidence both for segregated processing of color and form and also for a spatially distinct representation of conjunctions of these features. In support of a widely held view in physiology (Zeki 1978; Livingstone and Hubel 1984, 1988; DeYoe and Van Essen 1985; Ts'o and Gilbert 1988; Felleman and Van Essen 1991), we showed that the most informative voxels coding color were spatially segregated from those coding for orientation. We also showed that BOLD signals within hV4 produced a classification bias toward color decoding, supporting a role for hV4 in color processing (Zeki 1980; Vaina 1994; Bartels and Zeki 2000; Brewer et al. 2005). However, the main finding from this study was that, as early as primary visual cortex, patterns of BOLD signals distributed across voxels not only contain reliable information about orientation and color but also about the specific conjunction of these features. This suggests that a representation of color–orientation conjunctions may exist as early as V1, arising from activity that is spatially independent from activity associated with processing each feature separately.

Evidence That Conjunctions Are Coded throughout Visual Cortex

We obtained direct functional evidence for conjunction coding of color and orientation in the human visual brain. We found that information related to a specific pairing of color and orientation could be decoded from BOLD responses as early as primary visual cortex. As our analysis was able to distinguish between 2 double-conjunction stimuli that contained the same basic color and orientation information, BOLD activation patterns across voxels relating to either of these 2 stimuli could not have been informative to a classifier unless there were specific voxels tuned to specific feature conjunctions. The finding that conjunction information is coded in voxel ensembles that are spatially distinct from those coding the independent features comprising conjunctions supports the view that separate functional units may underlie the coding of conjunction information (Lennie 1998).

Though it is tempting to infer from the results of this study that specific conjunction detectors may exist early in human visual cortex, it is important to note that any inference made at a neuronal level is limited by the indirect relationship between

BOLD signals and neural activity (Logothetis et al. 2001). Furthermore, we cannot fully exclude the possibility that a classifier may have used a mixing of color-specific with orientation-specific neural responses along with a neurovascular nonlinearity to distinguish the 2 double-conjunction stimuli. One unlikely but possible scenario for this could involve the presence of voxels containing, for example, primarily separate R and CW detectors. Due to our alternation of stimuli during the double-conjunction conditions, responses of such voxels could theoretically saturate during 1 of the 2 stimuli. That this actually happened is unlikely, because in the above scenario, those voxels informing particularly well on single features would also inform particularly well on conjunctions, although we found the opposite to be the case. Nonetheless, it shows that like most fMRI approaches (including those using adaptation paradigms), this study is confined to conclusions at the level of voxels, not that of neurons. However, if saturation effects (of deoxygenation or blood flow) lead a voxel to differentiate conjunctions, this information may also be available to neural codes, as both oxygen abundance and blood flow can influence neural activity (Moore and Cao 2008). Because converging evidence from electrophysiology shows that neurons coding multiple features appear to form a special class of neurons within visual cortex that may be preferentially located in feedback layers within the cortical sheet, in our view, such neural responses most likely account for the results obtained in this study (Bartels 2009; Shipp et al. 2009). Future electrophysiological recordings will better distinguish whether conjunctions as studied here are coded in single neurons, cell assemblies, or inhomogeneous feedback signals.

Finally, the results of this human fMRI study provide evidence for a representation of conjunction information as early as primary visual cortex. As we designed each stimulus to drive all early cone-output channels equally (i.e., both stimuli contained isoluminant red and green spirals matched for color-channel load), it is unlikely that the decoding success of color–orientation conjunctions reflects an early segregation of orientation information defined by a parvocellular or koniocellular input (Sumner et al. 2008). Nevertheless, the findings of this study do not disentangle whether the conjunction information in V1 corresponds to the sensory encoding of specific color–orientation pairings or the perceptual readout of a "binding" operation. The use of stimuli that induce binding errors such as illusory conjunctions (Treisman and Schmidt 1982) may in future help to tease apart this mechanism.

References

Abbott LF, Varela JA, Sen K, Nelson SB. 1997. Synaptic depression and cortical gain control. Science. 275:220-224.

Barnes J, Howard RJ, Senior C, Brammer M, Bullmore ET, Simmons A, David AS. 1999. The functional anatomy of the McCollough contingent colour after-effect. Neuroreport. 10:195-199.

Bartels A. 2009. Visual perception: converging mechanisms of attention, binding, and segmentation? Curr Biol. 19:R300-R302.

Bartels A, Logothetis NK, Moutoussis K. 2008. fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. Trends Neurosci. 31:444-453.

Bartels A, Zeki S. 2000. The architecture of the colour centre in the human visual brain: new results and a review. Eur J Neurosci. 12:172-193.

Billock VA. 1995. Cortical simple cells can extract achromatic information from the multiplexed chromatic and achromatic signals in the parvocellular pathway. Vision Res. 35:2359-2369.

Billock VA, Tsou BH. 2004. A role for cortical crosstalk in the binding problem: stimulus-driven correlations that link color, form, and motion. J Cogn Neurosci. 16:1036-1048.

Brewer AA, Liu J, Wade AR, Wandell BA. 2005. Visual field maps and stimulus selectivity in human ventral occipital cortex. Nat Neurosci. 8:1102-1109.

Burges CJC. 1998. A tutorial on support vector machines for pattern recognition. Data Mining Knowledge Discovery. 2:121-167.

Clifford CW, Spehar B, Solomon SG, Martin PR, Zaidi Q. 2003. Interactions between color and luminance in the perception of orientation. J vision. 3:106-115.

Clifford CW. 2010. Dynamics of visual feature binding. In: Nijhawan R, Khurana B, editors. Space & time in perception & action. Cambridge (MA): Cambridge University Press. p. 199-215.

Clifford CW, Webster MA, Stanley GB, Stocker AA, Kohn A, Sharpee TO, Schwartz O. 2007. Visual adaptation: neural, psychological and computational aspects. Vision Res. 47:3125-3131.

Conway BR. 2001. Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). J Neurosci. 21:2768-2783.

Conway BR, Livingstone MS. 2006. Spatial and temporal properties of cone signals in alert macaque primary visual cortex. J Neurosci. 26:10826-10846.

Conway BR, Moeller S, Tsao DY. 2007. Specialized color modules in macaque extrastriate cortex. Neuron. 56:560-573.

Crowder NA, Price NS, Hietanen MA, Dreher B, Clifford CW, Ibbotson MR. 2006. Relationship between contrast adaptation and orientation tuning in V1 and V2 of cat visual cortex. J Neurophysiol. 95:271-283.

Derrington AM, Krauskopf J, Lennie P. 1984. Chromatic mechanisms in lateral geniculate nucleus of macaque. J Physiol. 357:241-265.

DeYoe EA, Van Essen DC. 1985. Segregation of efferent connections and receptive field properties in visual area 2 of the macaque. Nature. 317:58-61.

Engel SA. 2005. Adaptation of oriented and unoriented color-selective neurons in human visual areas. Neuron. 45:613-623.

Felleman DJ, Van Essen DC. 1991. Distributed hierarchical processing in the primate cerebral cortex. Cereb Cortex. 1:1-47.

Forte JD, Clifford CW. 2005. Inter-ocular transfer of the tilt illusion shows that monocular orientation mechanisms are colour selective. Vision Res. 45:2715-2721.

Friston KJ, Holmes AP, Poline JB, Grasby PJ, Williams SC, Frackowiak RS, Turner R. 1995. Analysis of fMRI time-series revisited. Neuroimage. 2:45-53.

Furmanski CS, Engel SA. 2000. An oblique effect in human primary visual cortex. Nat Neurosci. 3:535-536.

Gegenfurtner KR, Kiper DC, Levitt JB. 1997. Functional properties of neurons in macaque area V3. J Neurophysiol. 77:1906-1923.

Haynes JD, Rees G. 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nat Neurosci. 8:686-691.

Holcombe AO, Cavanagh P. 2001. Early binding of feature pairs for visual perception. Nat Neurosci. 4:127-128.

Hosoya T, Baccus SA, Meister M. 2005. Dynamic predictive coding by the retina. Nature. 436:71-77.

Hubel DH, Livingstone MS. 1990. Color and contrast sensitivity in the lateral geniculate body and primary visual cortex of the macaque monkey. J Neurosci. 10:2223-2237.

Johnson EN, Hawken MJ, Shapley R. 2001. The spatial transformation of color in the primary visual cortex of the macaque monkey. Nat Neurosci. 4:409-416.

Johnson EN, Hawken MJ, Shapley R. 2008. The orientation selectivity of color-responsive neurons in macaque V1. J Neurosci. 28:8096-8106.

Kaiser PK. 1991. Flicker as a function of wavelength and heterochromatic flicker photometry. In: Kulikowski JJ, Walsh V, Murray IJ, editors. Limits of vision. Basingstoke (United Kingdom): MacMillan. p. 171-190.

Kamitani Y, Tong F. 2005. Decoding the visual and subjective contents of the human brain. Nat Neurosci. 8:679-685.

Kiper DC, Fenstemaker SB, Gegenfurtner KR. 1997. Chromatic properties of neurons in macaque area V2. Vis Neurosci. 14:1061-1072.

Kohn A, Movshon JA. 2003. Neuronal adaptation to visual motion in area MT of the macaque. Neuron. 39:681-691.

Krekelberg B, Boynton GM, van Wezel RJ. 2006. Adaptation: from single cells to BOLD signals. Trends Neurosci. 29:250-256.

Kriegeskorte N, Goebel R. 2001. An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical mr volumes. Neuroimage. 14:329-346.

Lennie P. 1998. Single units and visual cortical organization. Perception. 27:889-935.

Liu J, Wandell BA. 2005. Specializations for chromatic and temporal signals in human visual cortex. J Neurosci. 25:3459-3468.

Livingstone MS, Hubel DH. 1984. Anatomy and physiology of a color system in the primate visual cortex. J Neurosci. 4:309-356.

Livingstone MS, Hubel DH. 1988. Segregation of form, color, movement, and depth: anatomy, physiology, and perception. Science. 240:740-749.

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature. 412:150-157.

Mannion DJ, McDonald JS, Clifford CWG. 2009. Discrimination of the local orientation structure of spiral glass patterns early in human visual cortex. Neuroimage. 46:511-515.

McCollough C. 1965. Color adaptation of edge-detectors in the human visual system. Science. 149:1115-1116.

Moore CI, Cao R. 2008. The hemo-neural hypothesis: on the role of blood flow in information processing. J Neurophysiol. 99:2035-2047.

Peirce JW, Solomon SG, Forte JD, Lennie P. 2008. Cortical representation of color is binocular. J vision. 8(6):1-10.

Sapountzis P, Schluppeck D, Bowtell RW, Peirce JW. 2010. A comparison of fMRI adaptation and multi-variate pattern classification analysis in visual cortex. Neuroimage. 49:1632-1640.

Sasaki Y, Rajimehr R, Kim BW, Ekstrom LB, Vanduffel W, Tootell RB. 2006. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. Neuron. 51:661-670.

Sawamura H, Orban GA, Vogels R. 2006. Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. Neuron. 49:307-318.

Sereno MI, McDonald CT, Allman JM. 1994. Analysis of retinotopic maps in extrastriate cortex. Cereb Cortex. 4:601-620.

Seymour K, Clifford CW, Logothetis NK, Bartels A. 2009. The coding of color, motion, and their conjunction in the human visual cortex. Curr Biol. 19:177-183.

Shipp S, Adams DL, Moutoussis K, Zeki S. 2009. Feature binding in the feedback layers of area V2. Cereb Cortex. 19:2230-2239.

Sumner P, Anderson EJ, Sylvester R, Haynes JD, Rees G. 2008. Combined orientation and colour information in human V1 for both L-M and S-cone chromatic axes. Neuroimage. 39:814-824.

Thorell LG, De Valois RL, Albrecht DG. 1984. Spatial mapping of monkey V1 cells with pure color and luminance stimuli. Vision Res. 24:751-769.

Tolias AS, Keliris GA, Smirnakis SM, Logothetis NK. 2005. Neurons in macaque area V4 acquire directional tuning after adaptation to motion stimuli. Nat Neurosci. 8:591-593.

Treisman A, Schmidt H. 1982. Illusory conjunctions in the perception of objects. Cognitive Psychol. 14:107-141.

Ts'o DY, Gilbert CD. 1988. The organization of chromatic and spatial interactions in the primate striate cortex. J Neurosci. 8:1712-1727.

Vaina LM. 1994. Functional segregation of color and motion processing in the human visual cortex: clinical evidence. Cereb Cortex. 4:555-572.

Vapnik VN. 1998. Statistical learning theory. New York: Wiley.

Wandell BA, Chial S, Backus BT. 2000. Visualization and measurement of the cortical surface. J Cogn Neurosci. 12:739-752.

Yoshioka T, Dow BM. 1996. Color, orientation and cytochrome oxidase reactivity in areas V1, V2 and V4 of macaque monkey visual cortex. Behav Brain Res. 76:71-88.

Zeki SM. 1978. Functional specialization in the visual cortex of the monkey. Nature. 274:423-428.

Zeki SM. 1980. The representation of colours in the cerebral cortex. Nature. 284:412-418.