University of Groningen

Multivoxel fMRI analysis of color tuning in human primary visual cortex

Parkes, Laura M.; Marsman, Jan-Bernard C.; Oxley, David C.; Goulermas, John Y.; Wuerger, Sophie M.

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Multivoxel fMRI analysis of color tuning in human primary visual cortex

Laura M. Parkes

Imaging Sciences and Biomedical Engineering, School of Cancer and Imaging Sciences, University of Manchester, UK

Jan-Bernard C. Marsman

Laboratory for Experimental Ophthalmology and BCN Neuroimaging Center, School of Behavioral and Cognitive Neurosciences, University Medical Center Groningen, University of Groningen, The Netherlands

David C. Oxley

Department of Physics, University of Liverpool, UK

John Y. Goulermas

Department of Electrical Engineering, University of Liverpool, UK

Sophie M. Wuerger

School of Psychology, University of Liverpool, UK

We use multivoxel pattern analysis (MVPA) to study the spatial clustering of color-selective neurons in the human brain. Our main objective was to investigate whether MVPA reveals the spatial arrangements of color-selective neurons in human primary visual cortex (V1). We measured the distributed fMRI activation patterns for different color stimuli (Experiment 1: cardinal colors (to which the LGN is known to be tuned), Experiment 2: perceptual hues) in V1. Our two main findings were that (i) cone-opponent cardinal color modulations produce highly reproducible patterns of activity in V1, but these were not unique to each color. This suggests that V1 neurons with tuning characteristics similar to those found in LGN are not spatially clustered. (ii) Unique activation patterns for perceptual hues in V1 support current evidence for a spatially clustered hue map. We believe that our work is the first to show evidence of spatial clustering of neurons with similar color preferences in human V1.

Introduction

Color processing begins in the retina with three cone types broadly tuned to either short (S), medium (M), or long (L) wavelengths of light. Retinal ganglion cells compare the outputs of the cone cells and this cone opponency is inherited by neurons in the Lateral Geniculate Nucleus (LGN), which then project to neurons in the primary visual cortex (V1). The mapping from cone cells to the LGN is well understood; much less is known about the transformation that occurs between the LGN and V1 and the spatial organization of color-tuned neurons in V1, hence the focus of the current work.

The LGN consists of magno-, parvo-, and koniocellular layers which are thought to map onto the 'luminance' ('L+M'), the 'red–green' ('L−M'), and the 'yellow–blue' channel ('S−(L+M)'), respectively (Casagrande, 1994; Chatterjee & Callaway, 2002; Chichilnisky & Baylor, 1999; Derrington, Krauskopf, & Lennie, 1984; De Valois, Abramov, & Jacobs, 1966; De Valois, Cottaris, Elfar, Mahon, & Wilson, 2000; Gegenfurtner & Kiper, 2003; Solomon, White, & Martin, 1999), termed 'cardinal directions' by Krauskopf, Williams, and Heeley (1982). Importantly, these cardinal directions do not map onto the perceptual hues of red, yellow, green, and blue (De Valois, De Valois, & Mahon, 2000; De Valois, De Valois, Switkes, & Mahon, 1997; Wuerger, Atkinson, & Cropper, 2005) but appear as red ('L−M'), cyan ('M−L'), lime ('(L+M)−S'), and violet ('S−(L+M)').

The presence of wavelength selective cells in V1 has been known for many years (Conway & Tsao, 2006; Tootell, Nelissen, Vanduffel, & Orban, 2004; Zeki, 1983), but the distribution of peak tuning in V1 is much flatter than in the LGN, implying that a substantial number of cortical neurons are tuned to intermediate color directions (Hanazawa, Komatsu, & Murakami, 2000; Lennie, Krauskopf, & Sclar, 1990; Wachtler, Sejnowski, &

Albright, 2003). Also, whereas neurons with similar color preferences are spatially clustered and arranged in distinct layers in LGN, much less is known about the spatial organization of color-tuned neurons in V1, with the exception of a recent optical imaging study in the macaque providing evidence for a spatially organized map of hue tuning (Xiao, Casti, Xiao, & Kaplan, 2007). We hypothesize a shift of tuning in V1 away from cone-opponent cardinal colors toward intermediate, non-opponent, color modulations, and that cells with similar tuning will be topographically arranged.

fMRI has been used to study color processing in V1 of humans. However, the amplitude of the blood oxygenation-level-dependent (BOLD) fMRI response in V1 has been shown to depend on both the color and contrast of the stimulus, making any inferences on color tuning difficult. Different cardinal color directions produce different BOLD amplitudes even with matched psychophysical detection thresholds (Mullen, Dumoulin, McMahon, de Zubicaray, & Hess, 2007) or equal cone contrasts (Engel, Zhang, & Wandell, 1997; Liu & Wandell, 2005). The suggestion of spatial clustering of neurons tuned to particular hues (Xiao et al., 2007) allows a new approach to study color tuning with fMRI, namely a multivoxel pattern analysis (MVPA) approach. Two studies of orientation tuning in V1 (Haynes & Rees, 2005; Kamitani & Tong, 2005) used such an approach to infer spatial information beyond traditional fMRI resolution by pooling information from weak orientation-selective features across a region. The fMRI response pattern across the region was shown to be unique for each orientation of visual grating presented.

In this work we apply the MVPA method to test our hypotheses. Namely, we aim to determine the degree of spatial clustering of neurons with similar color preferences in human V1.

Methods

Participants

Five healthy subjects (2 female, age 21–31 years) with normal or corrected-to-normal vision gave written informed consent to take part in this study. All participants had normal color vision as assessed with the Cambridge Color Test (Regan, Reffin, & Mollon, 1994). The study was approved by the Sefton Liverpool Research Ethics Committee.

Experimental design and stimuli

We used high resolution fMRI (1.5-mm in-plane resolution) to determine the spatial clustering of V1 neurons with similar color preferences using two different sets of color modulations: (i) 3 cardinal directions and (ii) 4 perceptual hues. Each subject was scanned on two separate occasions (Experiments 1 and 2). In each experiment, isoluminant color stimuli were presented for 12 s followed by 12 s of an isoluminant gray screen. The color stimuli had similar spatial and temporal parameters as the stimuli used by Liu and Wandell (2005); they consisted of flickering radial sinusoidal gratings (0.8 cycles/deg; 1.5 Hz; 20 degrees of visual angle diameter), reversing contrast between: the endpoints of the cardinal directions in Experiment 1 and the perceptual hues and isoluminant gray in Experiment 2 (cf. Figure 1). All stimuli were presented on a neutral gray background. To control for attention subjects were asked to perform a forced choice task throughout the experiment where they had to decide if the fixation shape was a circle or a square.

Experiment 1

In Experiment 1, the three cardinal color directions were used for the stimuli, each presented three times in a random order per run (216 s). There were six runs giving 18 presentations for each of the three colors. Each run had a different random order of color presentation (Figure 2).

The CIE xy coordinates of these three color modulations were given as follows (see also Figure 1A; labeled as 'Card IsolumRed', etc.): The red–green direction varied from $x = 0.358$, $y = 0.340$ (reddish) to $x = 0.267$, $y = 0.386$ (greenish); the lime–violet direction from $x = 0.391$, $y = 0.552$ (lime) to $x = 0.286$, $y = 0.289$ (violet). All color stimuli were presented on a gray background ($x = 0.313$, 0.362) of the same luminance (470 cd/m^2); all components of the stimulus were isoluminant with the gray background for the chromatic modulations. The third direction was achromatic (ranging from black to white) with the same chromaticities as the gray background. Based on the results of Liu and Wandell (2005) the cone contrasts (vector length in cone contrast space) for the red–green, lime–violet, and black–white modulations were chosen as 9%, 86%, and 60%, respectively, in order to give roughly similar BOLD responses in V1. Given the gamut of our LCD projector (dashed line in Figure 1A), the maximum available cone contrast modulation from isoluminant reddish to isoluminant greenish is only about 9% (cf. Figure 1).

Experiment 2

In a second experiment, on a different day, four perceptual hues were used, each presented three times in a random order per run (288 s). There were five runs giving 15 presentations for each of the four colors. Each run had a different random order of color presentation. For each observer the perceptual hues (i.e., red, green, yellow, blue) were assessed individually prior to the scan using a
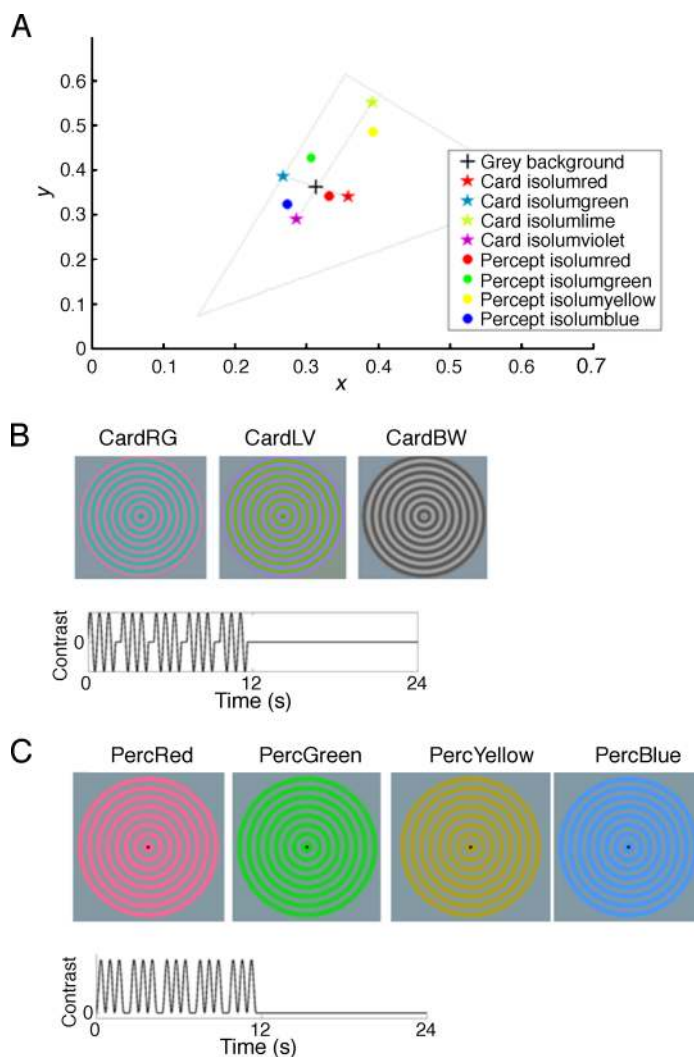
Figure 1. (A) The CIE *xy* chromaticity coordinates for all colors are shown. 'Card' refers to the cardinal color directions that optimally activate the layers of the LGN (Experiment 1) and 'Percept' refers to the perceptual hues (Experiment 2). (B) Spatial and temporal profiles of the three cardinal color modulations (Experiment 1) are shown (CardRG: modulations between cardinal red and green; CardLV: modulations between lime and violet; CardBW: modulations between black and white). The mean contrast is 0. (C). Spatial and temporal profiles of the four perceptual hue modulations (Experiment 2) are shown. The color varies over time between gray and the peak contrast of a particular perceptual hue.

modified hue-cancellation task (Wuerger et al., 2005). The contrast of the perceptual hues was then scaled such that they were roughly equal in terms of detectability (Webster, Miyahara, Malkoc, & Raker, 2000). We chose to use stimuli that were equally detectable rather than equidistant in the *xy* diagram, since CIE *xy* space is not uniform and distances in the *xy* diagram do not correspond to perceptual color differences. The average CIE coordinates were given as follows (see also Figure 1A): red: $x = 0.3312$,

$y = 0.3420$; green: $x = 0.3060$, $y = 0.4276$; yellow: $x = 0.3924$, $y = 0.4862$; blue: $x = 0.2730$, 0.3236. Stimuli were now modulated between the gray background color and a particular perceptual hue (cf. Figure 1C), with the mean contrast level being half the peak contrast of a particular hue. The luminance level of the test patterns was always the same as the background at each point in space and time.

The peak cone contrasts for the four hues were given as follows: red: 4%; green: 18%; yellow: 56%; blue: 16%. We used unipolar modulations between the gray background and a particular perceptual hue, since the perceptual hues (as assessed by a cancellation method) do not lie on a straight line including the gray background (Wuerger et al., 2005). All perceptual hues had the same luminance as the gray background (470 cd/m^2 for 2 subjects and 100 cd/m^2 for 3 subjects). Four subjects performed the 'cardinal colors' session first and 1 subject the 'perceptual hues' session first.

In the first session a retinotopic mapping scan was included, consisting of a rotating 45 degree checkerboard wedge, flickering at 8 Hz, with the same visual angle as the color stimuli. An inner radius of 0.75 degrees of visual angle was omitted from the retinotopic stimulus, thus avoiding the foveal tritanopic area of approximately 0.35 degrees in which S-cones are absent (Curcio et al., 1991). The wedge rotated a full circle in 48 s; there were 10 rotations.

Visual display system

Experiments were run on a standard DELL PC with a VSG2/5 graphics card (32-MB memory, Cambridge Research Systems). Stimulus presentation was controlled with Matlab 7 (Mathworks®) and stimuli were presented on a PANASONIC LCD PT-L785U projector, which was calibrated using a spectroradiometer (Photo Research PR650).
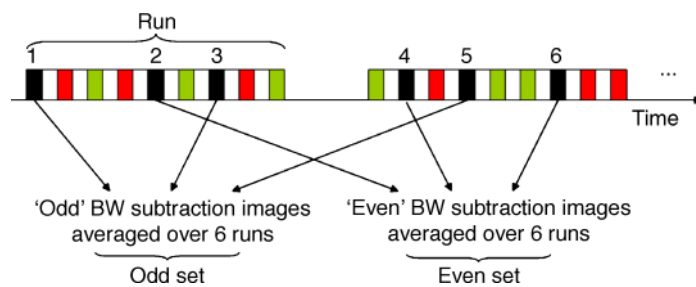

Figure 2. Paradigm. Each run consisted of nine 24-s presentations (12-s color stimulus, 12-s gray screen), hence lasting 216 s in total. The colors are presented in a random order over a total of 6 runs. For each color, the presentations are separated into two equal sets, named 'odd' and 'even' as shown.

fMRI acquisition

Scanning was performed on a 3 T Siemens system with an eight-channel phased-array head coil for signal collection. A high resolution EPI sequence with prospective motion correction (Thesen, Heid, Mueller, & Schad, 2000) was used for the color runs with the following parameters: TR 3 s, TE 35 ms, matrix 128×128, in-plane resolution 1.5 mm, slice thickness 2 mm, 29 slices covering the occipital cortex. In both scanning sessions a 1-mm isotropic 3D MPRAGE structural image was collected.

Analysis

The fMRI data were analyzed using BrainVoyager version 1.9 (Brain Innovations, Maastricht, The Netherlands) on an individual basis. Pre-processing comprised linear trend removal, motion correction, and slice time correction. No spatial or temporal smoothing was applied to the data from the color runs. The functional data were aligned onto the 3D anatomical image using coordinates contained in the image header files. The two anatomical images, one from each imaging session, were aligned automatically within BrainVoyager. Each alignment step was checked by eye and slight adjustments were made where necessary. A general linear model was created for each color run, comprising of either the three (cardinal colors) or four (perceptual hues) color regressors.

Identifying primary visual cortex (V1)

Standard procedures within BrainVoyager were used to segment white and gray matter and so identify the cortical boundary. The boundary is then inflated to give a smooth representation of the cortical surface. The functional data from the retinotopic run were analyzed using a box-car function convolved with a standard hemodynamic response function (HRF), for a range of 6 lag times (from 0 to 24 s). The lag time represents the visual angle of the stimulus, with 24 s covering one half of the visual field. The other half of the visual field was revealed by using an HRF shifted by 24 s. The significantly activated voxels were color-coded for lag time (visual angle) and displayed on the inflated cortical surface (see Figure 3B). Using knowledge of the retinotopic organization of the visual cortex (Wandell, Brewer, & Dougherty, 2005) the V1 region was identified by hand. For each subject, a binary mask of V1 voxels was produced, as shown in Figure 3.

We also created four additional mask images of sub-regions of V1: upper and lower visual fields based on the retinotopic maps and 'foveal' and 'eccentric' regions based on a crude division of V1 along the anterior–posterior axis.

Within V1 the average BOLD signal time course was plotted for each color stimulus for each subject. Baseline was set as the time period from 3 s before color stimulus onset to 3-s post-stimulus onset.

Creating subtraction images for each trial

The fMRI data from the color runs and the binary mask files were input to in-house software. For each presentation, 3 EPI images from 6 s to 15 s following color stimulus onset were averaged (i.e., those containing the maximum BOLD response), as were 2 EPI images from 3-s preceding color onset to 3-s post-onset (i.e., those containing the minimum BOLD response). These time periods were chosen by inspecting the average time course of the BOLD response in V1, avoiding the transition images between maximum and minimum responses. The average minimum image was subtracted from the average maximum image to give a single subtraction image per color presentation.

The stimuli were designed to give similar average BOLD response within V1. However, to avoid the classification results being driven by any small difference in overall amplitude rather than the spatial pattern of the response, the subtraction images were normalized within the region of interest (Haxby et al., 2001). The mean signal in the region of interest was found for each subtraction image and this was subtracted from each voxel value such that the average signal in the region of interest is now zero. To test this normalization process, for one subject, the signal in V1 was artificially elevated for one of the perceptual hues, by multiplying the voxel values by a
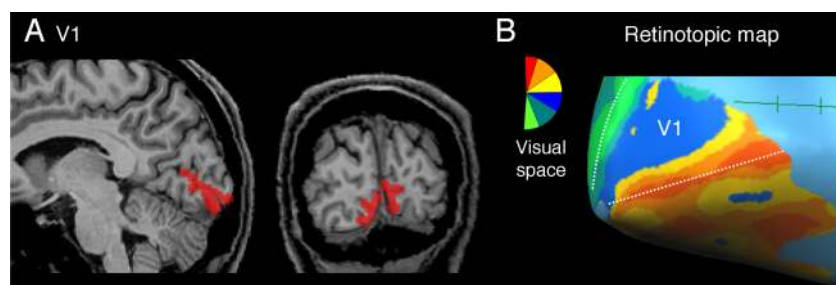

Figure 3. Mask regions. The figure shows mask regions of V1 (A) and the retinotopic map (B) in a typical subject. The color key shows that there are 6 discreet colors each coding for a different region of space.

factor from 1 to 10. The correlation and classification results were found in each case to determine if there was any effect of this global contrast enhancement.

Correlation analysis

To determine if there is a robust and unique pattern of fMRI activation for each color presented we looked at correlations between the activation patterns within V1, following a similar procedure to that described by Haxby et al. (2001). For each subject and each color category the subtraction images (see previous section) were split into two sets, namely even and odd sets, and averaged to create two mean subtraction images per color category (Figure 2).

We then determined whether the color the subject was viewing could be identified on the basis of the correlations between even and odd sets. Correlation coefficients were computed between the response patterns for odd sets compared to even sets for both same color (e.g., red odd sets and red even sets) and different colors (e.g., red odd and blue even). For the different-color case there were four possible combinations (red (odd or even) + blue (odd or even)), the mean of which was found. This was done on an individual basis and also on concatenated data from all subjects.

To determine the predictive power of the correlation data we asked the following question: Given the even data sets, does the highest correlation with a given odd data set correctly predict the color that was viewed? To calculate the confidence interval on a null result we used a bootstrapping technique and randomly permuted the correlation values 1000 times to find a mean performance and 95% confidence interval (1.96 * SD).

Pattern classification

For the classification we experimented with various machine learning algorithms (Duda, Hart, & Stork, 2001; Hastie, Tibshirani, & Friedman, 2001). As expected, due to the particular characteristics of the data sets (i.e., very small number of samples per class, with very large dimensionalities) complex classifiers, such as neural networks, did not perform robustly. The experiments showed that the specific data sets supported simpler classification surfaces that discriminate the different pairs of classes. Similar observations for such small sample cases have also been observed in other applications (Goulermas, Findlow, Nester, Howard, & Bowker, 2005). Here, we obtained the best classification performance using Regularized Discriminant Analysis (RDA; Friedman, 1989), a mixture between linear and quadratic discriminant analysis, where the class covariance estimates are replaced with regularized ones. Support Vector Machines (SVM) with a linear kernel (Joachims, 1999) gave similar results.

The set of subtraction images and the mask files were input into the classification software. Pair-wise classification was performed between all color pairs within each set, i.e., three pairings for the cardinal colors and six pairings for the perceptual hues. Two-way classification was used (rather than 3- or 4-way) in order to produce the largest number of tests to give the lowest confidence margin. Testing involved Leave-One-Out (LOO) validation, which is a special case of the k-fold cross validation widely used for classification analysis and classifier evaluation (Duda et al., 2001; Hastie et al., 2001). This is mandatory due to the finiteness of the data sets, to ensure that the classifier is trained with some samples, and tested with a completely different sample instances later, so that the final error assessment is objective. Given n data samples, LOO trains the classifier with $n - 1$ samples, and then tests with the remaining one. Subsequently, it repeats the process for a total of n times by alternating the left-out testing sample. Each testing prediction is recorded, and the final classification performance is the average result from all n runs.

To test whether the result is significantly above chance the procedure was repeated for a set of 'shuffled' data. The color classes were randomly shuffled 100 times in order to produce 100 sets of 'null' data where the performance should be at chance level. The same classification procedure as for the real data was repeated on these 100 sets. The mean and standard deviation were found and the 95% confidence interval (1.96 * SD) on this chance performance. If the real result lies outside the 95% confidence interval then it is deemed to be significant at the $p < 0.05$ level.

Results: Experiment 1

The V1 response to cardinal colors was tested in Experiment 1. Observers were presented with color modulations along the cardinal axes while fMRI data were obtained. Mean voxel numbers in V1 across all 5 subjects were 1850 with a standard deviation of 100.

Figure 4 shows the average BOLD response for the cardinal colors in V1, which were found to be similar for all colors.

Figure 5A shows fMRI activation patterns within V1 for each cardinal color in a typical subject. It is useful to inspect these raw data (an approach also followed by Haxby et al. (2001)) to assess whether the response patterns differ between conditions. In this case the response patterns for the odd and even sets are similar, showing that the response pattern is robust. However, the pattern for each color also appears similar. Figure 5B shows the correlation over voxels between two sets of the same color (RG–RG) and between two sets of different colors (RG–BW) for this subject. The correlations appear similar. Figure 5C shows the correlation coefficients between pairs of response patterns over all 5 subjects. On an individual basis, the correlation data could be used
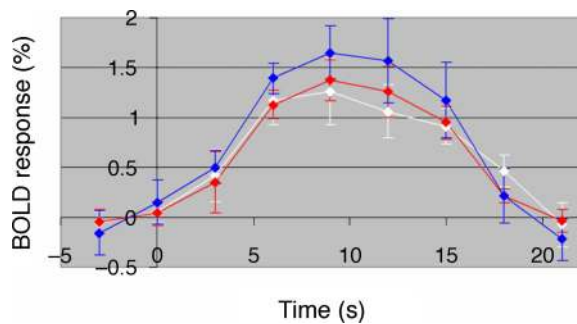
Figure 4. The BOLD response to cardinal colors in V1. The percentage BOLD increase is shown for all color stimuli, averaged over all subjects. The error bars show the *SD* over subjects. red = red–green, blue = lime–violet, white = black–white.
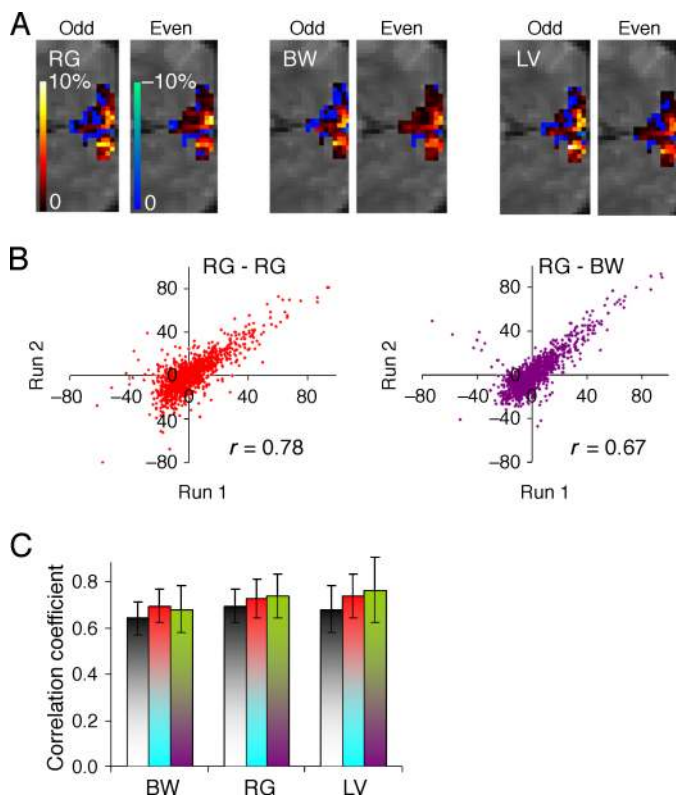

Figure 5. Correlation results for cardinal colors. (A) The BOLD response pattern within V1 is shown for one typical subject for even and odd sets of each of the color stimuli. The signal has been normalized by subtracting the mean signal over all voxels. Hence the colors represent the normalized BOLD signal amplitude with blue lower than mean and red/yellow higher than mean. The color bar indicates percentage BOLD signal change. (B) The lowest 'within' and 'between' category correlations for this subject. (C) Correlation coefficients concatenated across all subjects. The *x*-axis describes the color category for one data set and the column color represents the category of the second data set. Red/green is represented by the two-tone colors red/cyan and lime/violet is represented by lime/pink. The error bars represent the standard error over individual results from all subjects.
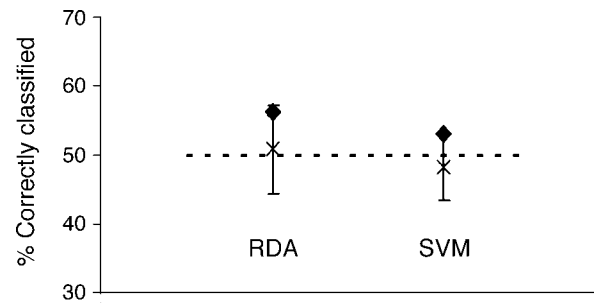

Figure 6. Classification results for cardinal colors. The figure shows the overall performance for all subjects and all pairs of color classifications for both the RDA and SVM analysis. The performance of the real data (black diamonds) is shown in comparison to the null result from 100 sets of shuffled data (black cross). The error bars represent the 95% confidence interval on the null result.

to correctly predict the color viewed in 47% of cases, which is within the 95% confidence interval for chance performance (33% ± 24%) and hence not significant. Hence the correlation data can provide no predictive power as to which color is viewed.

It is unsurprising therefore that the classification performances (Figure 6) are within the 95% confidence interval for chance performance and hence are not deemed to be significant.

Results: Experiment 2

In Experiment 1 we demonstrated that, although V1 neurons clearly respond strongly to cone-opponent cardinal color modulations, the color viewed could not be predicted based on the spatial distribution of activity. The main purpose of the second experiment was therefore to investigate whether spatial clustering of neurons with similar color preferences could be revealed with a different set of color stimuli. From neurophysiological data we know that a large proportion of neurons are tuned to hues different from the cardinal colors and a significant proportion of neurons responds best to a narrow range of hues and not to any linear combinations of cone outputs (Hanazawa et al., 2000; Lennie et al., 1990; Wachtler et al., 2003). An obvious candidate set of colors are the focal colors (red, green, yellow, blue) since not only are these important color categories to encode the statistical properties of our environment (Philipona & O'Regan, 2006), but neurons in areas as early as primary visual cortex may be particularly sensitive to these focal hues (De Valois, Cottaris et al., 2000; Horwitz, Chichilnisky, & Albright, 2007; Vautin & Dow, 1985). With this in mind, we measured the BOLD responses associated with the focal colors (red, green, yellow, blue), which will be referred to as the 'perceptual hues'.
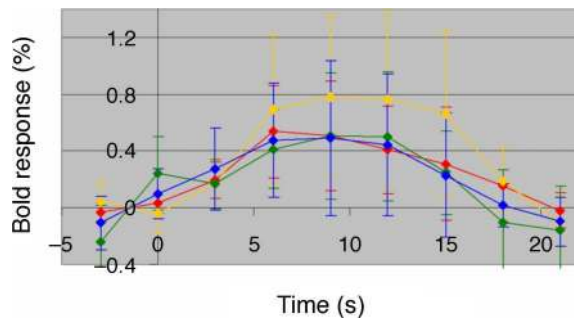
Figure 7. The BOLD response to perceptual hues in V1. The percentage BOLD increase is shown for all color stimuli, averaged over all subjects. The error bars show the *SD* over subjects. Line colors correspond to stimuli colors.

Figure 7 shows the average BOLD response for perceptual hues in V1. The responses are similar across all perceptual hues but lower than for cardinal colors (Figure 4). One explanation for the lower BOLD response for perceptual hues is that unipolar modulations were used rather than bipolar modulations and fewer neurons may respond to unipolar modulations or they may respond less strongly. Furthermore, the BOLD signal is contrast dependent and there is no obvious contrast metric one can use to equate different colors.

Figure 8A shows fMRI activation patterns within V1 for each perceptual hue in a typical subject. The response patterns for even and odd sets are similar, showing again that the response pattern is robust. Note that the individual patterns appear idiosyncratic and unique across different regions within V1; i.e., the differences do not appear to be driven by more global, larger scale differences in spatial pattern. Figure 8B shows the correlation over voxels between two sets of the same color (red–red) and between two sets of different colors (red–blue) for this subject. The correlations are different, with the same-color correlation being higher than the different color, showing that the pattern for each perceptual hue is unique. Figure 8C shows the correlation coefficients between pairs of response patterns over all 5 subjects, showing that this difference between 'same' and 'different' color correlations is maintained. It can be seen that the same-category correlations are generally higher than the different-category correlations. On an individual basis the correlation data could be used to correctly predict the color viewed in 55% of the cases, which is outside the 95% confidence interval on chance performance (25% ± 19%). Hence the color viewed in Experiment 2 can be predicted significantly above chance level from the fMRI data in V1.

Differences in same-category correlations (for example red–red is lower than blue–blue—Figure 8) may contribute to differences in different-category correlations and hence to the predictive power of the data, i.e., if one category has higher noise than another, the prediction may be made using this information. To assess this we
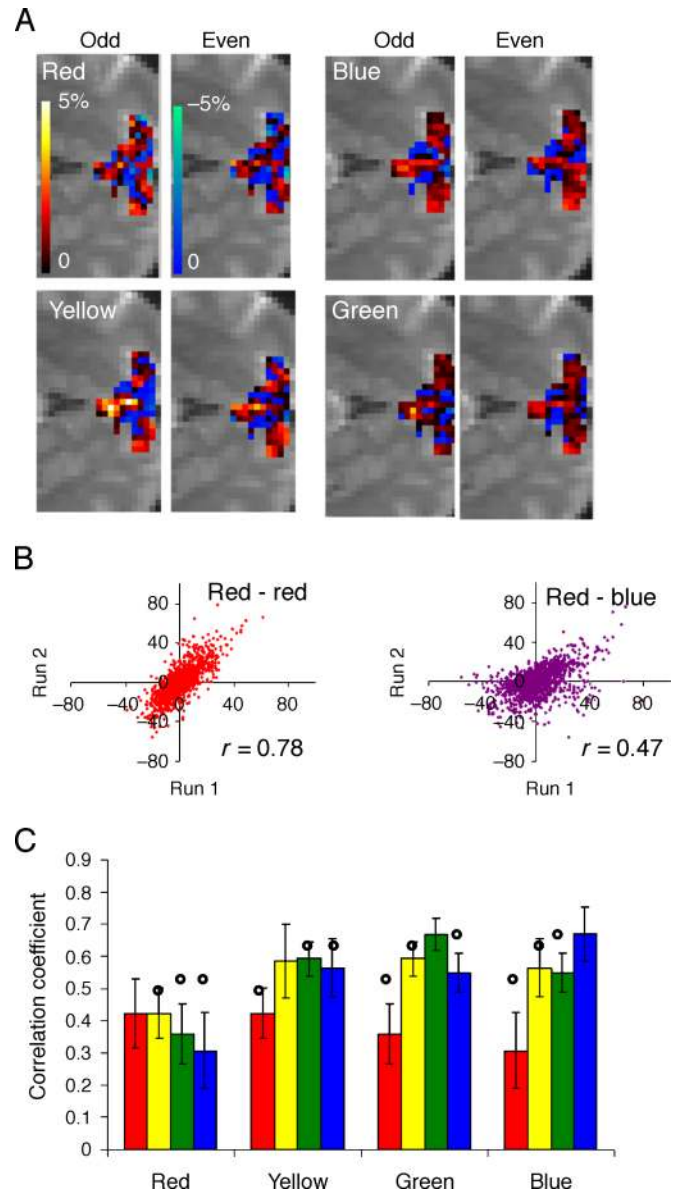

Figure 8. Correlation results for perceptual hues. (A) The BOLD response pattern within V1 is shown for one typical subject for the even and odd sets of each of the color stimuli. The signal has been normalized by subtracting the mean signal over all voxels. Hence the colors represent the normalized BOLD signal amplitude with blue lower than mean and red/yellow higher than mean. The color bar indicates percentage BOLD signal change. (B) Lowest 'within'- and 'between'-category correlations for this subject. (C) Correlation coefficients concatenated across all subjects. The x-axis describes the color category for one data set and the column color represents the category of the second data set. The error bars represent the standard error over individual results from all subjects. The circles show the predicted different-color correlation coefficients for the null hypothesis (see text).

considered the null hypothesis that all data have the same underlying response pattern but with different levels of noise. In this case, the different-category correlation coefficient is equal to the square root of the product of the same-category correlation coefficients (i.e., $r_{12} = \sqrt{(r_1 r_2)}$, see Appendix A). Using this we computed the different-category correlation coefficients for the null hypothesis for individual data and also the group result (Figure 8C) based on same-category coefficients only. We found that the predictive power of this simulated data was reduced to 25%, i.e., chance performance. Hence we conclude that the predictive power of 55% is due purely to the different underlying spatial patterns of the different categories.

Figure 9 shows the classification performance for both the RDA and SVM analysis. The perceptual hues are classified significantly above chance for both analyses.

There is a possibility that this above-chance classification performance is driven by differences in color preference over large regions, for example between upper and lower visual fields, rather than by local differences. To test for this we repeated the classification in sub-regions of V1: upper and lower visual fields and 'foveal' and 'eccentric' regions (see Methods section). Figure 10 shows that classification remained above chance in these sub-regions.

Figure 11 shows color-preference images (similar to the orientation-preference maps of Kamitani and Tong (2005), ocular dominance maps of Cheng, Waggoner, and Tanaka (2001), and category-preference maps of Grill-Spector, Sayres, and Ress (2006)) revealing the preference of each individual voxel to a particular color (i.e., the voxel is colored according to the stimulus to which it has the largest BOLD response). In general the images show scattered color preference, with no regional preference to a particular color, and are variable and idiosyncratic across subjects.

Figure 10. Classification results for perceptual hues within sub-regions of V1: eccentric (mean voxels = 420) and foveal (mean voxels = 980) regions of interest and upper (mean voxels = 1080) and lower (mean voxels = 780) visual fields. The figure shows the overall performance for all subjects and all pairs of color classifications using the RDA.

Our stimuli were chosen to give comparable BOLD responses (see Methods section) in order to give similar signal to noise ratios across stimuli within each color set. Figures 4 and 7 show that this was reasonably successful. However, our classification method does not make use of the average BOLD response in a region, but rather the *distribution* of responses within that region. Crucially, we normalized the average BOLD response within each region for each stimulus to remove this information from the classification procedure. This ensures that the difference in response distribution due to color information is
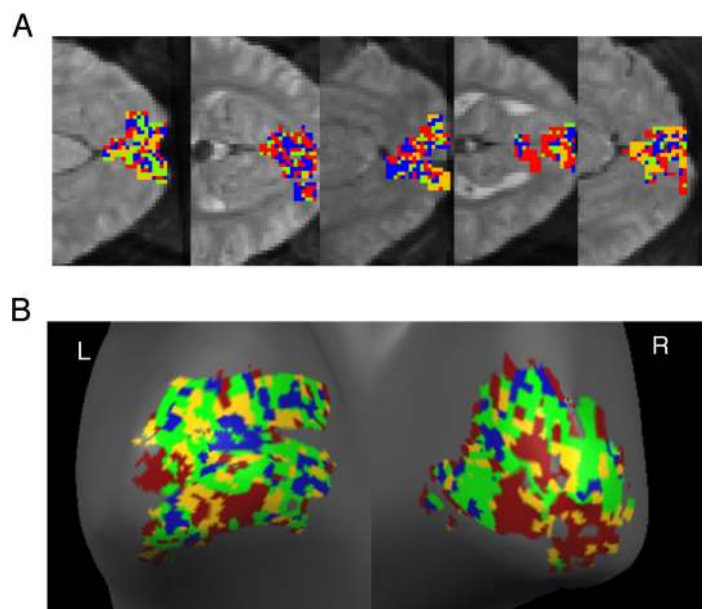
Figure 11. The 'color-preference' images show the preference of each voxel for a particular color. Each voxel is colored according to the perceptual hue (red, yellow, green, or blue) to which the voxel showed the largest BOLD response. (A) Results from all 5 subjects for the middle slice. (B) The results on the inflated cortical surface for one subject.
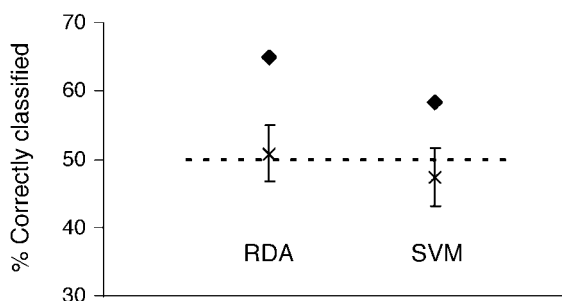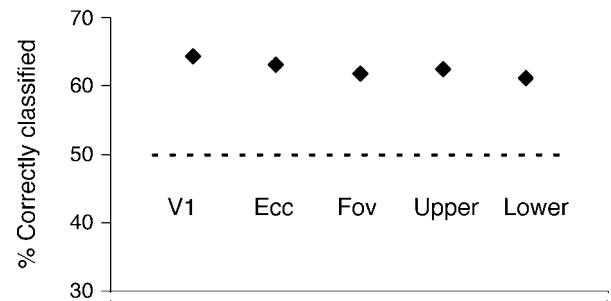

Figure 9. Classification results for perceptual hues. The figure shows the overall performance for all subjects and all pairs of color classifications for both the RDA and SVM analysis. The performance of the real data (black diamonds) is shown in comparison to the null result from 100 sets of shuffled data (black cross). The error bars represent the 95% confidence interval on the null result.

Figure 12. Effect of artificially increasing the global signal (by the contrast factor) for one perceptual hue on the SVM classification result in V1 for one subject.

used as opposed to contrast information (which would affect the global signal rather than the response pattern).

Figure 12 shows the success of this normalization procedure on data where the average BOLD response to one color was artificially increased. There is very little effect on the classification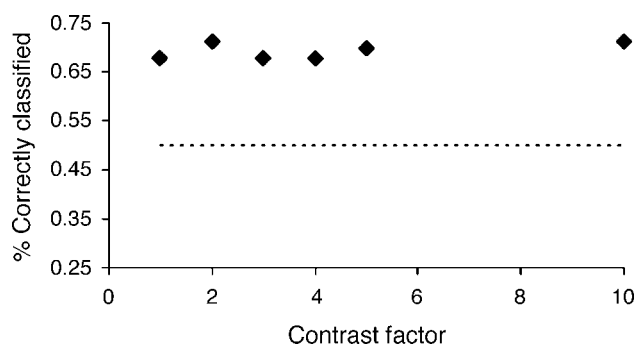 result even when considering a 10-fold increase (contrast factor = 10) in average response for one of the colors. We found no effect of this contrast manipulation on the correlation results.

Discussion

The results of Experiment 1 showed that the BOLD response to cardinal colors in V1 is high and the patterns of activation are robust (Figure 5). However, there is no significant difference between within- and between-category correlations, such that the color viewed could not be predicted (Figures 5C and 6). The most parsimonious explanation is that a significant number of V1 neurons respond to the cardinal colors, but that these neurons are not spatially clustered hence producing no unique response patterns. This explanation is supported by the large BOLD response for cardinal colors in V1 (Figure 4) and is in agreement with neurophysiological data in macaque showing that a significant number of V1 neurons are tuned to cardinal color directions (Hanazawa et al., 2000; Horwitz et al., 2007; Lennie et al., 1990; Wachtler et al., 2003).

The results of Experiment 2 showed that the BOLD response to perceptual hues in V1 is also high and the patterns of activation are robust (Figure 8). But most importantly, the within-category correlations are higher than the between-category correlations, such that the color viewed can be predicted significantly above chance (Figures 8C and 9). This suggests that the V1 neurons that respond to perceptual hues are spatially clustered, producing unique response patterns.

Spatial clustering of V1 neurons with similar chromatic preferences

The results of Experiment 2 show that perceptual hues generate unique response patterns (Figure 8), suggesting that the neurons that are responsive to perceptual hues are spatially clustered. In V1, cone-opponent cardinal color modulations generate a strong BOLD signal but poor correlation/classification results (Experiment 1); this lack of within-category correlation implies that neurons that preferentially respond to these color modulations are not spatially clustered in V1. Unipolar perceptual hue modulations, on the other hand, yield strong within-color correlations that are predictive of the color viewed. This suggests that perceptual hues are associated with a unique activation pattern, which in turn requires some spatial clustering of neurons with similar chromatic preferences.

Indeed, this highlights the key difference in the information we can get from our multivoxel analysis approach compared to previous work using univariate approaches. Previous fMRI work in this field has considered the mean BOLD response over a region to stimuli of different colors and contrasts. This allows the investigation of the overall chromatic sensitivity but does not provide direct information about the underlying chromatic tuning and spatial clustering of cortical neurons. Classification/correlation techniques may be a more sensitive technique to study chromatic tuning as they make use of the overall activity patterns across many voxels and are to a large extent independent of the overall amplitude. This is advantageous since the overall amplitude of the BOLD response to a particular color depends on both the hue and the contrast of the color signal.

The classification results alone do not however inform us of the nature of the spatial organization. It is possible that the results are driven by global rather than local differences—for example upper and lower visual fields may show slight differences in sensitivity to different colors. We addressed this issue by calculating classification performance for upper and lower visual fields alone and found that classification remained above chance and only marginally lower than when considering the entire V1 region (Figure 10). There is also recent evidence for changes in color preference with eccentricity (Mullen et al., 2007; Vanni, Henriksson, Viikari, & James, 2006). It was more difficult for us to test any potential effect of this on classification performance as we did not collect eccentricity maps, however, a crude segmentation of V1 into smaller 'foveal' and more 'eccentric' regions again revealed little reduction in classification performance (Figure 10). In addition the 'color-preference' map (Figure 11) shows no regional bias to one particular color compared to another. The activation patterns for each hue appear to vary locally rather than globally.

It may have been expected that cardinal colors would show significant classification due to the known differences in contrast sensitivity for each color with increasing eccentricity (Mullen et al., 2007; Mullen & Kingdom, 2002). However, while the behaviorally assessed contrast sensitivity is robust and consistent across subjects (Mullen & Kingdom, 2002), the fMRI response is variable (Mullen et al., 2007), which could explain why we do not see this effect. In addition, our stimuli were at a higher contrast level than in the previous studies, which could reduce any difference between the stimuli. There is little behavioral evidence for any difference in response to unipolar perceptual hues with increasing eccentricity (Parry, McKeefry, & Murray, 2006), hence it seems unlikely that this could contribute to the significant classification results that we see in this case.

Methodological issues

One important issue is the limit of spatial information that the MVPA technique can access. Clearly, as the size of tuning column becomes smaller and the fMRI voxel larger, the bias between voxels for a particular color will become progressively smaller. Optical imaging data in the macaque (Xiao et al., 2007) suggest that color columns are on the order of 150 μm apart, i.e., ten-fold smaller than the 1.5 mm voxel size used in our study. Earlier work using MVPA to successfully classify orientation information used voxel sizes of 3 mm (Haynes & Rees, 2005; Kamitani & Tong, 2005). Orientation columns are known to have a width of 300–500 μm in the macaque (Vanduffel, Tootell, Schoups, & Orban, 2002); i.e., a similar ratio of column/voxel size as for our study. However, the limit of detectability also depends on the signal to noise ratio, which is dependent on the square of the voxel size and hence 4-fold smaller for our study. This reduction in power could contribute to the lower classification performance that we find in comparison to the earlier work on orientation tuning.

In general we found the correlation results to be more successful in predicting the color viewed than the classification results. This could be due to the fact that, for the correlation analysis, data are averaged over half of the total number of presentations, hence boosting the signal to noise ratio. The classification approach instead considers individual subtraction images. Both approaches can give information on underlying spatial organization of neurons responding to different colors.

The classification/correlation results are driven by the spatial pattern of response as the normalization of the data ensures that the overall amplitude is not driving the classification. Figure 12 confirms the success of this normalization approach, showing that, even if the overall response to one color is amplified by a factor of 10, the classification result is not significantly affected. The correlation results are also unaffected. The difference between the within- and between-category correlations, which drive the color predictions, are therefore due to unique spatial patterns of BOLD response associated with each color. A number of alternative normalization approaches could have been considered. Our subtraction approach is similar to that used by Haxby et al. (2001), whereas others divided by the mean signal (Haynes & Rees, 2005), and others chose not to normalize (Eger, Ashburner, Haynes, Dolan, & Rees, 2008). Our subtraction method will remove the mean signal difference between conditions but there could be residual differences in variance. However, this is likely to be small in our case due to the similar BOLD responses across color stimuli (Figures 4 and 7).

It is possible that the poor classification performance for cardinal colors is due to saturation of the BOLD signal (i.e., blood vessels are maximally expanded), which is indeed larger for cardinal colors (Figure 4) than for perceptual hues (Figure 7). However we feel that this is unlikely as the 1–1.5% BOLD signal we report is not as high as typical BOLD signal amplitudes seen at the same field strength in V1. For example, recent work by Leontiev and Buxton (2007) shows a 1.9% BOLD response to a checkerboard stimulus and 4% BOLD response to hypercapnia in V1. Hence, it seems unlikely that the BOLD response has reached its maximum possible value in our study. It is however possible that the neuronal response is saturated due to the high color contrasts used. Even if the neuronal response is saturated, so long as the BOLD response is not saturated, this should not affect our classification performance. The classification relies on biases between voxels (the activation *pattern*) due to differences in populations of neurons with similar chromatic preferences, not due to the different absolute BOLD response to different colors.

In Experiment 1 we used cone-opponent color modulations. Under these conditions we could not identify a unique activation pattern. We think the most parsimonious explanation is that the neurons tuned to these cone-opponent modulations are not spatially clustered. This lack of unique activation pattern could be due to the opponency of the bipolar modulations and not due to the particular colors tested. However the aim was to test whether similar spatial clustering of these particular cone-opponent neurons as seen in the LGN also exists in V1. Our data suggest that this is not the case. In Experiment 2 we employed perceptual hues, namely modulations between gray and a particular hue. We chose the focal hues of red, green, yellow, and blue because of their behavioral importance (Philipona & O'Regan, 2006) and neurophysiological evidence suggesting increased sensitivity of V1 neurons to these colors (De Valois, Cottaris et al., 2000). We were able to reveal spatial clustering of neurons with similar chromatic preference, but this does not imply that this clustering is exclusively for focal colors, and it is the topic of further work to investigate whether hue maps can be revealed with intermediate

colors. Since we used unipolar modulations in Experiment 2 and bipolar modulations in Experiment 1, we cannot directly compare these two experiments to draw conclusions about the precise chromatic tuning of V1 neurons. Such inferences are beyond the scope of this study: the purpose of these experiments was to test whether multivoxel pattern analysis is a sensitive tool to reveal spatial clustering of neurons with similar chromatic preferences.

Why does a spatial hue map make sense?

From a functional point of view, in analogy with orientation discrimination, an orderly columnar organization of hues provides the most parsimonious model to explain one major function of color vision, namely discriminating visual stimuli on the basis of hue. Optical imaging data from the macaque show evidence for such an orderly hue map in V1 (Xiao et al., 2007). If such a spatial hue map exists in primary visual cortex, we would also predict that colors that are perceptually more similar should show stronger correlations than colors that are perceptually dissimilar. Our experiment was not designed to test this particular prediction, but the correlation results for perceptual hues (Figure 8C) are at least not inconsistent with such a prediction, e.g., response pattern to 'green' has a higher correlation with the response pattern to 'yellow' and 'blue' than to 'red'. This suggests that neurons responding to green are physically closer in space to those tuned to yellow or blue than to those tuned to red, which is consistent with recent findings using optical imaging (Chen & Zhu, 2001; Miki, Liu, & Liu, 2004). At this stage we can only speculate whether the predicted correlation patterns would arise when tested with more appropriate stimuli. The current study was not designed to correlate color similarity with classification/correlation performance, but it demonstrates that multivoxel correlation analysis of BOLD responses may be a useful tool to identify brain areas that correspond to perceptual measures, such as color similarity.

In summary, we have shown that the multivoxel fMRI analysis can be used to reveal the spatial clustering of neurons with similar chromatic preferences. In V1, neurons with cone-opponent properties (cardinal colors) were found not to be spatially clustered, but we found evidence for a spatial hue map consistent with recent optical imaging data.

Appendix A

Assuming that each color produces the same underlying fMRI response pattern, but with different noise levels, we can calculate the predicted different-color correlations based on the same-color correlations.

Consider the same-color correlation coefficient for color X, defined as

$$\rho_{X_1X_2} \equiv \frac{\sigma^2_{X_1X_2}}{\sqrt{\sigma^2_{X_1} \cdot \sigma^2_{X_2}}} = \frac{\sum(x_1 - \overline{x}_1)(x_2 - \overline{x}_2)}{\sqrt{\sum(x_1 - \overline{x}_1)^2 \sum(x_2 - \overline{x}_2)^2}}, \quad (A1)$$

i.e., the X_1 and X_2 covariance divided by the product of the standard deviations. We can assume that the standard deviation of X_1 is the same as for X_2 as these are created from two groups of data drawn from the same underlying population, giving

$$\rho_{X_1X_2} = \frac{\sigma^2_{X_1X_2}}{\sigma^2_{X_1}}. \quad (A2)$$

Consider a second color Y, which has the same underlying pattern as color X but with added noise. Assuming that the noise is independent to the signal, the addition of noise to X will not affect the covariance only the standard deviations. Therefore,

$$\sigma^2_{Y_1Y_2} = \sigma^2_{X_1X_2} = \sigma^2_{X_1Y_2} = A, \quad (A3)$$

where A is the true signal covariance. Hence the between-color correlation is given by

$$\rho_{X_1Y_2} \equiv \frac{\sigma^2_{X_1Y_2}}{\sqrt{\sigma^2_{X_1} \cdot \sigma^2_{Y_2}}} = \frac{A}{\sqrt{\left(\frac{A}{\rho_{X_1X_2}}\right)\left(\frac{A}{\rho_{Y_1Y_2}}\right)}} = \sqrt{\rho_{X_1X_2}\rho_{Y_1Y_2}}. \quad (A4)$$

This was also numerically verified by simulation.

This treatment assumes that the same-color correlation coefficients are noise free, and uses them to drive the predictions. We thank one of the reviewers of this manuscript for suggesting a more accurate approach of treating the four same-color correlation coefficients as parameters of a model, and ask how well the model can fit the entire set of 16 correlation coefficients. This then recalculates the same-color correlation coefficients according to the data from the entire data set. These values are then used to compute the different-color correlation coefficients as described above. Again, we found the predictive power of this simulated data was only 25%, i.e., chance performance.

Corresponding author: Laura M. Parkes.
Email: Laura.Parkes@manchester.ac.uk.
Address: Imaging Science and Biomedical Engineering, School of Cancer and Imaging Sciences, Stopford Building, The University of Manchester, Oxford Road, Manchester M13 9PT, UK.

References

Casagrande, V. A. (1994). A third parallel visual pathway to primate area V1. *Trends in Neuroscience, 17,* 305–310. [PubMed]

Chatterjee, S., & Callaway, E. M. (2002). S cone contributions to the magnocellular visual pathway in macaque monkey. *Neuron, 35,* 1135–1146. [PubMed] [Article]

Chen, W., & Zhu, X. H. (2001). Correlation of activation sizes between lateral geniculate nucleus and primary visual cortex in humans. *Magnetic Resonance in Medicine, 45,* 202–205. [PubMed]

Cheng, K., Waggoner, R. A., & Tanaka, K. (2001). Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron, 32,* 359–374. [PubMed] [Article]

Chichilnisky, E. J., & Baylor, D. A. (1999). Receptive-field microstructure of blue–yellow ganglion cells in primate retina. *Nature Neuroscience, 2,* 889–893. [PubMed]

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

Curcio, C. A., Allen, K. A., Sloan, K. R., Lerea, C. L., Hurley, J. B., Klock, I. B., et al. (1991). Distribution and morphology of human cone photoreceptors stained with anti-blue opsin. *Journal of Comparative Neurology, 312,* 610–624. [PubMed]

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. L., Abramov, I., & Jacobs, G. H. (1966). Analysis of response patterns of LGN cells. *Journal of the Optical Society of America, 56,* 966–977. [PubMed]

De Valois, R. L., Cottaris, N. P., Elfar, S. D., Mahon, L. E., & Wilson, J. A. (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]

De Valois, R. L., De Valois, K. K., & Mahon, L. E. (2000). Contribution of S opponent cells to color appearance. *Proceedings of the National Academy of Sciences of the United States of America, 97,* 512–517. [PubMed] [Article]

De Valois, R. L., De Valois, K. K., Switkes, E., & Mahon, L. (1997). Hue scaling of isoluminant and cone-specific lights. *Vision Research, 37,* 885–897. [PubMed]

Duda, R. O., Hart, P. E., & Stork, D. G. (2001). *Pattern classification* (2nd ed.). New York: John Wiley and Sons.

Eger, E., Ashburner, J., Haynes, J. D., Dolan, R. J., & Rees, G. (2008). fMRI activity patterns in human LOC carry information about object exemplars within category. *Journal of Cognitive Neuroscience, 20,* 356–370. [PubMed]

Engel, S., Zhang, X., & Wandell, B. (1997). Colour tuning in human visual cortex measured with functional magnetic resonance imaging. *Nature, 388,* 68–71. [PubMed]

Friedman, J. H. (1989). Regularized discriminant-analysis. *Journal of the American Statistical Association, 84,* 165–175.

Gegenfurtner, K. R., & Kiper, D. C. (2003). Color vision. *Annual Review of Neuroscience, 26,* 181–206. [PubMed]

Goulermas, J. Y., Findlow, A. H., Nester, C. J., Howard, D., & Bowker, P. (2005). Automated design of robust discriminant analysis classifier for foot pressure lesions using kinematic data. *IEEE Transactions on Biomedical Engineering, 52,* 1549–1562. [PubMed]

Grill-Spector, K., Sayres, R., & Ress, D. (2006). High-resolution imaging reveals highly selective nonface clusters in the fusiform face area. *Nature Neuroscience, 9,* 1177–1185. [PubMed]

Hanazawa, A., Komatsu, H., & Murakami, I. (2000). Neural selectivity for hue and saturation of colour in the primary visual cortex of the monkey. *European Journal of Neuroscience, 12,* 1753–1763. [PubMed]

Hastie, T., Tibshirani, R., & Friedman, J. (2001). *The elements of statistical learning*. New York: Springer-Verlag.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in

ventral temporal cortex. *Science, 293,* 2425–2430. [PubMed]

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

Horwitz, G. D., Chichilnisky, E. J., & Albright, T. D. (2007). Cone inputs to simple and complex cells in V1 of awake macaque. *Journal of Neurophysiology, 97,* 3070–3081. [PubMed] [Article]

Joachims, T. (1999). Making large-scale SVM learning practical. In B. Schölkopf, C. Burges, & A. Smola (Eds.), *Advances in kernel methods: Support vector learning.* Cambridge, USA: MIT Press.

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

Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research, 22,* 1123–1131. [PubMed]

Lennie, P., Krauskopf, J., & Sclar, G. (1990). Chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience, 10,* 649–669. [PubMed] [Article]

Leontiev, O., & Buxton, R. B. (2007). Reproducibility of BOLD, perfusion, and CMRO2 measurements with calibrated-BOLD fMRI. *Neuroimage, 35,* 175–184. [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]

Miki, A., Liu, C. S., & Liu, G. T. (2004). Effects of voxel size on detection of lateral geniculate nucleus activation in functional magnetic resonance imaging. *Japanese Journal of Ophthalmology, 48,* 558–564. [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]

Mullen, K. T., & Kingdom, F. A. (2002). Differential distributions of red–green and blue–yellow cone opponency across the visual field. *Visual Neuroscience, 19,* 109–118. [PubMed]

Parry, N. R., McKeefry, D. J., & Murray, I. J. (2006). Variant and invariant color perception in the near peripheral retina. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 23,* 1586–1597. [PubMed]

Philipona, D. L., & O'Regan, J. K. (2006). Color naming, unique hues, and hue cancellation predicted from

singularities in reflection properties. *Visual Neuroscience, 23,* 331–339. [PubMed]

Regan, B. C., Reffin, J. P., & Mollon, J. D. (1994). Luminance noise and the rapid-determination of discrimination ellipses in color deficiency. *Vision Research, 34,* 1279–1299. [PubMed]

Solomon, S. G., White, A. J., & Martin, P. R. (1999). Temporal contrast sensitivity in the lateral geniculate nucleus of a New World monkey, the marmoset Callithrix jacchus. *The Journal of Physiology, 517,* 907–917. [PubMed] [Article]

Thesen, S., Heid, O., Mueller, E., & Schad, L. R. (2000). Prospective acquisition correction for head motion with image-based tracking for real-time fMRI. *Magnetic Resonance in Medicine, 44,* 457–465. [PubMed]

Tootell, R. B., Nelissen, K., Vanduffel, W., & Orban, G. A. (2004). Search for color 'center(s)' in macaque visual cortex. *Cerebral Cortex, 14,* 353–363. [PubMed] [Article]

Vanduffel, W., Tootell, R. B., Schoups, A. A., & Orban, G. A. (2002). The organization of orientation selectivity throughout macaque visual cortex. *Cerebral Cortex, 12,* 647–662. [PubMed] [Article]

Vanni, S., Henriksson, L., Viikari, M., & James, A. C. (2006). Retinotopic distribution of chromatic responses in human primary visual cortex. *European Journal of Neuroscience, 24,* 1821–1831. [PubMed]

Vautin, R. G., & Dow, B. M. (1985). Color cell groups in foveal striate cortex of the behaving macaque. *Journal of Neurophysiology, 54,* 273–292. [PubMed]

Wachtler, T., Sejnowski, T. J., & Albright, T. D. (2003). Representation of color stimuli in awake macaque primary visual cortex. *Neuron, 37,* 681–691. [PubMed] [Article]

Wandell, B. A., Brewer, A. A., & Dougherty, R. F. (2005). Visual field map clusters in human cortex. *Philosophical Transactions of the Royal Society B: Biological Sciences, 360,* 693–707. [PubMed] [Article]

Webster, M. A., Miyahara, E., Malkoc, G., & Raker, V. E. (2000). Variations in normal color vision. II. Unique hues. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 17,* 1545–1555. [PubMed]

Wuerger, S. M., Atkinson, P., & Cropper, S. (2005). The cone inputs to the unique-hue mechanisms. *Vision Research, 45,* 3210–3223. [PubMed]

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

Zeki, S. (1983). Colour coding in the cerebral cortex: The reaction of cells in monkey visual cortex to wavelengths and colours. *Neuroscience, 9,* 741–765. [PubMed]