

The uses of colour vision: behavioural and physiological distinctiveness of colour stimuli

Andrew M. Derrington*, Amanda Parker, Nick E. Barraclough, Alexander Easton, G. R. Goodson, Kris S. Parker, Chris J. Tinsley and Ben S. Webb

School of Psychology, University Park, Nottingham NG7 2RD, UK

Colour and greyscale (black and white) pictures look different to us, but it is not clear whether the difference in appearance is a consequence of the way our visual system uses colour signals or a by-product of our experience. In principle, colour images are qualitatively different from greyscale images because they make it possible to use different processing strategies. Colour signals provide important cues for segmenting the image into areas that represent different objects and for linking together areas that represent the same object. If this property of colour signals is exploited in visual processing we would expect colour stimuli to look different, as a class, from greyscale stimuli. We would also expect that adding colour signals to greyscale signals should change the way that those signals are processed. We have investigated these questions in behavioural and in physiological experiments.

We find that male marmosets (all of which are dichromats) rapidly learn to distinguish between colour and greyscale copies of the same images. The discrimination transfers to new image pairs, to new colours and to image pairs in which the colour and greyscale images are spatially different.

We find that, in a proportion of neurons recorded in the marmoset visual cortex, colour-shifts in opposite directions produce similar enhancements of the response to a luminance stimulus.

We conclude that colour is, both behaviourally and physiologically, a distinctive property of images.

Keywords: colour vision; image segmentation; object detection; figure-ground; monkey; vision

1. DISTINCTIVENESS OF COLOUR STIMULI

(a) *Background*

This paper is concerned with the fact that we readily distinguish between stimuli that contain colour variations and those that do not. Quite apart from any significance that we might attach to particular colours, or any use we may choose to make of colour information, we distinguish between the presence and the absence of colour in a scene, a pattern or a picture.

This distinctiveness of colour might be something that we learn from our experience of colour and greyscale images in photography, in the cinema and on television. However, this paper makes the case that the distinctiveness could arise naturally from the visual consequences of the presence of colour in an image. We shall argue in support of this latter position in three ways. First, we shall argue that the contribution of colour and, in particular, hue variations to image segmentation makes colour images qualitatively different from greyscale images. Second, we shall show behavioural evidence that a monkey, the common marmoset *Callithrix jacchus*, rapidly learns to discriminate between colour and greyscale images. Finally, we shall show physiological results indicating that the

presence of colour modulates the responses of some cortical neurons to a luminance stimulus. Crucially, in this last instance, the nature of the modulation depends only on the presence of colour; it is independent of the actual colour that is present.

In the following two subsections, we consider the role of colour, and particularly of hue information, in the identification of stimuli and in the segmentation of images.

(i) *Hue and stimulus identification*

Specific hues, or hue combinations, may be used to detect and identify specific objects or classes of object. For example, many edible fruits may be detected against a background of foliage by the fact that they are somewhat redder than the foliage. This type of information, in which specific hues indicate specific types of object, is obviously important. Indeed, there is evidence to support the proposal that in old-world monkeys the photoreceptor pigments that support red–green colour discriminations have evolved to maximize the sensitivity of the discrimination between fruit and foliage (Osorio & Vorobyev 1996; Regan *et al.* 1998; Sumner & Mollon 2000). However, there is no reason to suppose that links between specific hues and specific objects would lead to a situation where colour images, as a class, look different from greyscale images. Each hue might have its own significance, but the importance of a coloured image would depend on which hues it contained, not on the mere presence of the variations in hue that differentiate colour from greyscale

* Author for correspondence (andrew.derrington@nottingham.ac.uk).

One contribution of 14 to a Discussion Meeting Issue 'The physiology of cognitive processes'.

images. We consider a more general kind of information provided by hue variations that could make their mere presence in an image important.

(ii) *Hue and image segmentation*

One of the fundamental problems of vision is image segmentation—deciding which bits of the image should be grouped together because they represent parts of the same object. Indeed, the segregation of objects from each other and from the background is a prerequisite for most of the cognitive operations that can be performed on visual stimuli. Edge detection, the detection of lines in the image along which the luminance changes abruptly, has always been regarded as an important early computation in image segmentation (Marr 1976, 1982; Marr & Hildreth 1980). However, although edges may indicate the boundaries between objects of different reflectance, luminance edges can also be misleading because they can be caused by shadows. This problem can be confounded when objects occlude parts of each other so that a single object may be divided into several distinct areas. Hue provides two types of information that can resolve these two kinds of confusion.

First, retinal images, like most images received from the natural world are formed by light reflected from a group of objects illuminated by a single light source. It is a physical fact that, in such images hue boundaries cannot be caused by shadows. Hue boundaries must represent boundaries between differently reflecting materials, and thus potentially between different objects. Consequently, algorithms for segmenting images can be driven by searching the image for boundaries between areas of different hue (Hurlbert 1989).

Second, even when they are in different levels of shadow, if the different parts of a partially occluded object are made of the same material and lit by the same light source, they will have the same hue, though not necessarily the same lightness. Thus, in principle, the different parts of an image that represent the same object can be linked together by visual processing, which links together areas of the same hue (Mollon 1989).

Figure 1 shows these aspects of colour in a diagram that represents the colour information available to the animal that we used in our behavioural experiments, the male marmoset. Male marmosets are dichromats, so the information available to them about any single colour can be represented in a two-dimensional plot showing the relative extent to which it excites both photoreceptor systems (Tovée *et al.* 1992; Travis *et al.* 1988). Thus, each colour can be represented as a point in this space. We are interested in colour boundaries, such as that between an object and its background. The convention that we adopt is to represent each boundary as an arrow with its tail on the point representing the background colour and its head on the point representing the foreground colour.

The thickly drawn arrows in figure 1 all point either towards or away from the origin. They represent colour changes that do not change the ratio of excitations of the two photoreceptors. They are changes of intensity without any change in hue: each arrow represents a boundary that could be caused by a shadow. In fact, the heavy arrows in figure 1 are all running along the same line pointing towards the origin, and thus all represent different inten-




Figure 1. Conventions for representing colour and luminance boundaries in images. Each point in the diagram represents the information about a hypothetical colour provided by the photoreceptors of a dichromat, such as the male marmoset: the excitations produced in the two photoreceptor systems, the L cones and the S cones. The L-cone excitation is plotted against the S-cone excitation. Borders between two areas of different colour within an image are represented as an arrow from one colour (usually the background colour) to the other (usually the foreground colour). The heavy arrows indicate boundaries that involve a change in lightness but no change of hue; they could be caused by shadows. The paler arrows indicate changes of hue (and often also lightness). In an image of reflecting objects lit by a single source, changes of hue indicate changes in the reflecting material.

sities of the same hue. Their different start- and endpoints could all represent the same object seen in different levels of shadow. The fact that they lie on the same line radiating from the origin could be used to link together different parts of the same object.

Different hues are encoded as different ratios of photoreceptor activation and would lie on lines radiating out from the origin in different directions. Shadows on objects of each hue would be represented by arrows pointing up or down the hue line, towards the origin or away from it. The thinly drawn arrows in figure 1 represent colour-shifts that could not be caused by shadows. The clue to this is that they point across the diagram along lines that do not pass through the origin. They represent boundaries in the image that are associated with changes in hue and thus necessarily correspond to changes in reflectance.

(b) *Testable predictions*

These two potential uses of hue in image segmentation mean that all images that contain hue variations form a distinct class for a visual system. Regardless of the actual hues involved, all images that contain more than one hue have the common property that colour algorithms can be used to segment them into patches corresponding to differently reflecting surfaces, and to link together the




Figure 2. Experimental procedure: photograph of a marmoset retrieving his reward after rotating the stimulus cube. See text for further details.

patches that correspond to the same surface in different levels of shadow. Thus, if colour information is used in segmentation, we can make two predictions.

First, colour and greyscale images should naturally form different categories and it should be easy to teach animals to respond differentially to them. Specifically, once trained to respond differentially to colour and greyscale images, they should extend the differential response—without further training—to novel pairs of colour and greyscale images, including pairs in which the colour image contains only hue variations not present in the training set. In § 2, the behavioural section of this paper, we confirm this prediction by showing that male marmosets trained to distinguish between pairs of greyscale and colour images transfer the discrimination to novel images and to novel colours.

Our second prediction is that, at some point in the visual pathway, hue boundaries should modulate visual responses in a distinctive way. The modulation should be different from that produced in neurons whose responses are used to identify hue, in which different hues produce different responses and complementary hues usually produce opposite responses. Rather, what should happen is that the presence of a hue boundary should modulate the response of the neuron in a way that does not depend on the actual hues. The neuron should signal the presence of the hue boundary rather than either the hues that are associated with it, or its sign. In § 3, the physiological section of this paper, we demonstrate that a small number of neurons in the striate and prestriate cortex have their responses modulated by hue boundaries in this way.

2. BEHAVIOURAL CLASSIFICATION OF COLOUR AND GREYSCALE STIMULI

(a) Behavioural training and testing

(i) General method

Male common marmosets (*Callithrix jacchus*) were trained to discriminate between pairs of images in a variant of the Wisconsin General Testing Apparatus (Dias *et al.* 1996a,b) that could be mounted on the side of their home cage (see figure 2). One image was the positive, or rewarded stimulus, S+, and the other, the unrewarded

stimulus, S-. Each image was printed on card and mounted behind a transparent cover on the front face of a backless hollow plastic cube, which measured *ca.* 25 mm on each side. Each cube was mounted on a vertical spindle and could be rotated by the marmoset in order to retrieve the reward, a small piece of marshmallow, from the back of the cube containing the positive stimulus.

Before the start of each trial, the marmoset was signalled to move to the back of his cage, 0.6 m from the stimulus cubes. Two screens, one opaque and one transparent, were placed in front of the stimulus cubes, which were then loaded with the stimuli and the reward for the next trial. To start the trial, the opaque screen was removed, whereupon the marmoset inspected the stimuli and signalled his choice by moving towards one of the stimulus cubes. His choice was recorded and the transparent screen was removed, allowing him access to his chosen stimulus cube. If he was correct, he was able to rotate the cube, retrieve the reward and consume it, before returning to the back of the cage for the next trial.

Using this procedure, marmosets performed between 40 and 100 trials per day while being fed a normal food ration. Feeding took place after testing each day.

(ii) Stimuli

Stimuli were photographic-quality prints 23 mm square, produced by an inkjet printer (Epson Stylus Photo 750). Training stimuli were clip-art bitmaps printed in colour or greyscale. Testing stimuli were complex two-dimensional shapes drawn by doodling with a Bezier line tool in a drawing package (see figure 4a-c) and printed as a solid shape of a single colour on a grey or coloured background.

The reflectance spectra of printed samples of the colours used in the test stimuli were measured with a Data-color SF600 reflectance spectrophotometer at 10 nm intervals between 360 and 700 nm, and interpolated at 4 nm intervals. The emission spectra of the lamps in the behavioural testing apparatus were measured at 4 nm intervals using a SpectraResearch PR650 spectroradiometer. The marmoset cone absorption spectra were calculated at 4 nm intervals using the polynomial approximation from Stavenga *et al.* (1993). Products of sets of emission, reflectance and absorption spectra were used to calculate the coordinates of the test-stimulus colours in a dichromatic cone-excitation space for the marmoset.

(iii) Pre-training and training

Pre-training was designed to teach the marmosets to select the rewarded stimulus. Once they had habituated to the apparatus, marmosets were trained to retrieve pieces of marshmallow from the stimulus cubes. Then they were trained to rotate the stimulus cubes to retrieve marshmallow. Finally, they were trained to move to the back of the cage at the start of a trial before beginning trials in which a pair of coloured clip-art pictures were used as stimuli and one of the images was consistently rewarded.

Ten pairs of pictures, chosen so that they contained a wide range of colours and image features, and so that no colour or image feature consistently predicted reward, were used as discrimination tests for pre-training. Monkeys were trained for 40 trials per day on each discrimination problem until they chose the positive stimulus on




Figure 3. (a) Number of errors made by four monkeys in learning a colour/greyscale image pair to a criterion of 10 consecutive correct choices, plotted against the number of discriminations learned. For Giles (squares) and Piers (circles) the colour image was rewarded; for Rupert (triangles) and Hugo (diamonds) the greyscale image was rewarded. After the sixth discrimination the monkeys were trained for 2 days on a list of the first five discriminations. (b) Performance of the same four monkeys learning a list of 10 novel colour/greyscale discriminations over four days. Each monkey was trained on the list for 40 trials per day.

at least 36 of the 40 trials on three consecutive days. Throughout training and testing, the correct stimulus was alternated from side to side using a Gellermann sequence.

Once pre-training with mixed stimuli had been completed, the marmosets were divided into two groups and began learning to discriminate between colour and greyscale versions of the same clip-art picture. For two of them (Giles and Piers), the colour picture was rewarded and for the other two (Rupert and Hugo), the greyscale picture was rewarded. Once established for a given monkey, the association between colour and reward was retained throughout his training. At this stage, we excluded any pictures containing red, so that we could use red pictures to test whether the colour discrimination transfers to novel colours. Each colour/greyscale discrimination was learned to a criterion of 10 consecutive correct responses.

Figure 3a shows the numbers of errors made by four monkeys learning to distinguish between colour and greyscale versions of the same picture to a criterion of ten consecutive correct responses. Performance on the first six discriminations was extremely variable. In the best case, Rupert had an immediate preference for the second greyscale image and so learned the correct response with no errors. In the worst case, Giles made almost 90 errors while learning the fourth image pair. All four learned the fifth image with no more than one error, but the sixth image produced variable performance.

At this point, they were retrained to ensure that their performance was 90% correct over 40 trials for the first five image pairs presented as a list, i.e. on each trial, any of the five image pairs could be presented with the constraint that no pair could be presented for the n th time

until all had been presented $n - 1$ times. Giles, Rupert and Piers reached 90% correct on the first day of training with the list, but Hugo took 2 days.

After this confirmation that they remembered the first five discriminations that they had learned, they resumed training with individual colour/greyscale discriminations. The last four discriminations were learned very rapidly. There were 20 errors by the four monkeys, and the modal number of errors per discrimination was zero. This suggests that, by this stage, the monkeys had learned to rely on the presence of colour to enable them to choose the correct stimulus and did not need to learn each image pair separately.

Figure 3b confirms this impression. It shows performance on a novel list of 10 colour/greyscale discriminations, which were presented in 40 trials per day over four days. Even on the first day, performance is between 80 and 95% correct. It improves slightly over the four days and is between 90 and 95% correct on day 4. This consistently high level of performance with novel stimuli indicates that, like humans, marmosets reliably distinguish images that contain colour variations from those that contain only variations in lightness: colour is a distinctive stimulus quality.

(iv) Necessary controls

In order to establish that the monkey's discrimination performance is based on the property of colour, we need to exclude two possibilities. First, we need to exclude the possibility that the marmosets are choosing between coloured and grey-scale images because of luminosity differences rather than colour differences. Although we have no




Figure 4. Stimuli for colour discrimination testing. (a) Training stimuli and test logic for testing discriminations without differentially rewarding the discriminanda. In 80% of the trials, the monkeys were trained for reward on the list of four colour/greyscale pairs. In 20% of the trials, a pair of test stimuli was presented and both stimulus boxes contained a reward. The monkeys' performance on test trials informed us about their stimulus choices in the absence of differential reward. (b) Five of the 10 test stimuli used in demonstrating that the colour/greyscale discrimination is not based on an apparent luminance or contrast difference. (c) Five of the 16 test stimuli used in demonstrating that the colour discrimination transfers to a novel hue. (d) Training- and test-stimulus colours plotted in cone-excitation coordinates with arrows indicating the borders between background and figure colours. The three heavy arrows show the figure/ground borders of the training stimuli. The lighter, numbered arrows show the figure/ground borders of the test stimuli. No. 1 is the blue test in (b). No. 2 is the red 'novel hue' test from (c), which was successfully discriminated by Hugo, Rupert and Piers but not by Giles. No. 3 is the lighter red that was successfully discriminated by Giles. Note that the heavy arrows representing the training stimuli all run across the diagram in a direction of increasing relative S-cone excitation (down and right), whereas the arrows representing the novel hue stimuli run up and left in a direction of increasing L-cone excitation.

a priori reason to suppose that there are luminance differences between colour and greyscale stimuli, we do not know the marmoset's spectral luminosity function and it is possible that the colours that we used might appear to be systematically brighter (or darker) than the greys with which they are matched. Second, we need to exclude the possibility that the marmosets have simply learned a list of colours.

(v) Test procedure

In order to measure how the training transferred to specific sets of test image pairs without differentially rewarding the responses to test image pairs and without losing the effect of training, we used the following procedure.

The animal was first trained on a list of four colour/greyscale image pairs to a criterion of 90% correct in the 40 trials of a day's testing. Then, during testing sessions, 80% of the trials were 'training' trials and would use a stimulus pair from the trained list, and the reward would only be received if the correct image was chosen (colour or greyscale in accordance with the training history of the particular monkey). Twenty per cent of trials, selected at random, were 'test' trials. On test trials, a test stimulus pair was presented and both of the stimulus cubes contained a reward. Consequently, the monkey's responses on test trials are determined exclusively by his training and not by reward contingencies on previous test trials.

(b) Control for luminosity differences

To eliminate the possibility that the marmosets use luminosity differences rather than colour differences, we prepared a series of colour images that consisted of different shapes, all of the same blue on the same grey background (see figure 4b). The comparison greyscale images

Table 1. Performance of the four monkeys on the different transfer tests.

test	Giles (%)	Hugo (%)	Rupert (%)	Piers (%)
luminance transfer	97.5	92.5	97.5	100
transfer to novel hue (dark red)	41	95	100	98
transfer to novel hue (light red)	97	—	—	—
dissimilar image pairs	80	80	92.5	85

used the same grey background but, in 50% of them, the shapes were lighter than the background and in the other 50% darker than the background. Test pairs of images drawn at random from this list were used for every fifth trial while the marmosets were performing discriminations for reward using a list of four greyscale and colour stimuli that they had previously learned (see figure 3a). Whenever the luminosity test pairs were presented, both choice boxes were loaded so that whichever choice they made would be rewarded. This 'double baiting' enables us to discover how the monkey's training controls his choice from each test pair without training them on the test pair (Walsh *et al.* 1993).

Table 1 shows the performance of each monkey on the transfer tests. The overwhelming majority of their choices were consistent with the colour discrimination (Giles 97.5%, Hugo 92.5%, Piers 100% and Rupert 97.5% on the luminance transfer test). Clearly, their colour choices are not based on luminosity differences between colours and greys because otherwise they would have made incorrect choices consistently, either on the light grey or on the dark grey images, depending on whether the blue that we had chosen appeared brighter or darker than the background.

(c) *Transfer to novel colours and novel image pairs*

We used a similar procedure to rule out the possibility that the marmosets are simply using a catalogue of colours, compiled during the experiment, to guide their choices. To do this, we made a set of 16 test image pairs in which the coloured images all used shapes made from a deep red on a grey background (see figure 4c). None of the previously used coloured images had contained any red. As before, test image pairs selected at random were presented on every fifth trial while the marmosets were performing discriminations for reward using a pre-learned list of four greyscale and colour stimuli. Both of the images of the colour test pair were rewarded. Hugo, Piers and Rupert transferred their colour discrimination to the novel red images: they scored 95%, 98% and 100%, respectively. Giles, however, did not transfer his colour discrimination to red images: his mean score was 41%.

One possible reason for Giles's failure is that he was unable to determine that the red test pattern was coloured because he could not distinguish between the red test colour and a dark grey. To investigate this possibility, we tried to train him to distinguish the red from dark grey by rewarding him only if he made the correct choice on trials with red test patterns. His mean performance in 160 trials over four days was 54% correct. Clearly, he was unable to make the discrimination, although his performance reached 58% on the fourth day, suggesting that he might

eventually have learned to do so. However, there can be no doubt that the reason why he did not transfer the colour discriminations to red images is that, during the time he was training with these, he was unable to distinguish the red from a dark grey. It is possible that Giles's failure to make this discrimination arises because he has a different colour vision phenotype from the others, but it is equally probable that it is a difficult discrimination for all of them. We have since learned, with other stimuli, that Giles is slower than the other three to learn difficult visual discriminations.

Having established that Giles's failure to generalize the colour/greyscale discrimination to red images was caused by his difficulty in discriminating that the stimulus was coloured, we decided to test him again using a lighter version of the same hue, which should be easier to discriminate. We changed the dark red to a lighter red (see figure 4d) and repeated the transfer test. He transferred almost perfectly: his average score was 97% correct over 8 days. Thus, despite Giles's early failure, all the monkeys transferred the colour/greyscale discrimination to one or other of the novel hues.

Finally, we used the same transfer procedure to test all the monkeys to see whether the colour/greyscale discrimination would generalize to a series of 10 greyscale/colour image pairs in which the greyscale and colour images were different spatial patterns. Although the transfer was not perfect, all of the monkeys scored at substantially above chance (Giles 32 out of 40, Hugo 32 out of 40, Piers 34 out of 40, Rupert 37 out of 40). The difference between the lowest of these scores and chance performance is extremely significant ($p < 0.005$; χ^2 -test).

These results show that, once they have solved a number of discriminations in which colour is the common factor, male marmosets rapidly and efficiently utilize colour to choose between new stimuli, even when those stimuli differ in other respects. They classify novel image pairs according to whether they are colour or greyscale, even when those images contain only hues that are not present in the training set. They even categorize pairs of pictures that differ in other respects, according to the colour/greyscale criterion.

The fact that an abstract image attribute, such as colour, is used in such a versatile and consistent way to guide the choice between novel stimuli is consistent with the notion that colour is a highly conspicuous stimulus property. The rapid transition from needing to learn each discrimination by reinforcement training to immediate, almost errorless performance with new discriminations (shown in figure 3a) suggests that the conspicuousness of colour is not simply a consequence of the reinforcement contingencies learned in our experiments, but is innate.

The fact that the colour/greyscale discrimination trans-

fers to novel colours on which the marmoset has not been trained (shown in table 1) is consistent with the notion that, like us, the marmoset includes all colour images in a common category and it is the selection of this category that has been reinforced. The most likely reason for the perceptual salience of stimuli that contain hue variations is the usefulness of hue variations in image segmentation. This leads to the prediction that monochrome pictures, which are coloured but contain only a single hue, should be classified as 'non-coloured' images. We have not tested this prediction because it is not possible to present coloured monochrome pictures on a coloured monochrome background in our present apparatus. We are at present building apparatus that will make it possible to present such stimuli in order to test this prediction.

The next question to consider is whether the behavioural distinctiveness of colour for the marmoset is reflected in the neurophysiology of cortical visual processing.

3. MODULATION OF PHYSIOLOGICAL RESPONSES BY COLOUR-SHIFTS

In considering what kind of physiological properties might support the behavioural discrimination between the presence and absence of colour, it is, first, important to distinguish it from normal colour discriminations. In normal colour discriminations, the aim is to identify the hue rather than merely to detect the presence or absence of variations in hue. Neurons that support hue identification must respond differentially to varying hue. For example, they may only respond to a narrow range of hues (Zeki 1983*a,b*) or they may respond in an opponent fashion to complementary hues (Wiesel & Hubel 1966). Neither of these types of response will support discrimination between the presence and absence of colour. An essential requirement is a response that signals the presence of hue variations in the stimulus but that does not differentiate between hues.

A second requirement arises if we assume that the functional significance of the ability to distinguish between the presence and the absence of colour is likely to be related to the use of hue in image segmentation. Hue boundaries allow us to identify a set of edges in the image that represent reflectance changes, and hence could not be caused by shadows. In general, these hue boundaries will coincide with luminance edges. Under this assumption, the appropriate neural signal would be a change in the response to a luminance stimulus when it is associated with a hue boundary, rather than an overt response to the hue boundary.

In order to look for neural signals that fulfil both of these requirements, we have measured the responses of neurons in marmoset visual cortex to small patches of moving sinusoidal grating. The grating was of optimal spatial frequency and orientation (i.e. the spatial frequency and orientation that produced the largest response), and the patch was of optimal length and width. The grating was presented under three different conditions (see figure 5). In the 'control' condition, the hue and mean luminance of the patch of grating was the same as the surrounding screen. In the two 'test' conditions, the patch was changed in hue either in a bluish or in a yellowish direc-




Figure 5. Stimuli for the physiological experiment, plotted in cone-excitation coordinates. The point in the centre of the plot represents the baseline colour and luminance of the display monitor. To its right and left are the points representing the colour backgrounds. The heavy arrows represent the range of colours covered by each of the sinusoidal gratings. The light arrows represent, for each background, the range of colour changes between the screen surrounding the grating patch (which remained at the baseline colour and luminance of the monitor) and the different parts of the grating. The different symbols represent the same information plotted for L-cones of the three different possible peak wavelengths. Circles represent 563 nm, inverted triangles 556 nm and upright triangles 545 nm.

tion while the surrounding screen was maintained at the same hue and luminance. Thus, at the boundaries of the grating patch, the screen changed both in luminance and in hue.

(a) Physiological recording

Physiological preparation for recording from neurons with microelectrodes and for presenting visual stimuli were essentially as have been described previously (Felisberti & Derrington 2001).

(i) Visual stimulation

Visual stimuli—spots or patches of different colours and luminances—and patches of sinusoidal grating were presented on a Sony Multiscan GDM200PDST monitor at a frame rate of 120 Hz by a Macintosh computer running custom-written software. Stimuli were adjusted in position and size and, if appropriate, in orientation and spatial frequency, in order to maximize the response of the neuron. Multiple stimuli, patches of different sizes, colours and contrasts, which could be combined with sinusoidal gratings of different contrasts, were randomly interleaved so that responses to stimuli that were being compared were accumulated concurrently. Stimulus luminances were set using a lookup table to compensate for the nonlinear relation between the luminance and applied voltage of the display. For coloured stimuli, cone exci-

tations were calculated using measurements of the emission spectrum of the display screen, made at 4 nm intervals in the range 380–780 nm. Cone spectra were calculated using the approximations of Stavenga *et al.* (1993) for peak wavelengths of 423 nm for the short-wavelength-sensitive cone, and 545, 556 and 563 nm for the different possible long-wavelength-sensitive cones.

The colour changes of the grating patch are illustrated in cone-excitation coordinates in figure 5. They were calculated to produce a large variation in the short-wavelength cone signal and minimal variation in the long wavelength cone signal. Such colour changes produce a strong response in colour-opponent neurons driven by blue cones. However, in this case we are not interested in direct responses to the changes in hue of the grating patch. Instead, we are interested in the possibility that the change in hue will modulate the response to luminance.

Figure 5 shows the hue boundaries associated with the gratings on patches of different hues. The heavy double-headed arrows connect the colours corresponding to the peaks and troughs of the sinusoidal grating. When the patch of grating is the same hue as its surround, the arrows are aligned along a line pointing to the origin. Both of the changes in hue of the grating patch change the lateral position of the arrow, without changing its orientation, so that it is no longer aligned with the origin. Thus, a mechanism designed to 'mark' spatial modulations of colour that include hue would mark the luminance changes within these gratings when they are presented on the coloured patches.

The thin arrows in figure 5 show the colour changes that occur at the edge of the patch of sinusoidal grating. They connect the point that represents the luminance and colour of the unmodulated display with the points that represent the luminances and colours of the peaks, the troughs and the zero crossings of the sinusoidal gratings on coloured patches. These colour changes also include hue shifts—none of the arrows lies on a line that passes through the origin. Thus, both at their edges and within the patches of sinusoidal grating, there are strong hue shifts.

(b) *Response of a single neuron to colour and luminance modulations*

Figure 6 shows how the hue signals modulate the responses of a striate cortical neuron to a sinusoidal grating. Histograms *a–c* show the response to the moving grating presented alone (*a*) and superimposed on the two coloured patches (*b, c*). When the grating is presented on a coloured patch, the response is bigger by *ca.* 20%. The two complementary coloured patches cause approximately equal enhancements of the response magnitude. However, when the patch colours are presented in isolation (histograms *e* and *f*), they elicit no response.

The coloured patches are spatially and temporally uniform: the colour-shift is turned on at the start of the stimulus presentation and off at the end. This raises the question of whether the lack of response in histograms *e* and *f* is because the neuron is insensitive to the colours of the patches, or because it is insensitive to their spatial and temporal frequencies.

Histograms *g* and *h* answer the question. They show the responses of the same neuron to colour modulations at




Figure 6. Responses of a neuron in striate cortex to the different stimuli illustrated in figure 5. (*a*) Response to a sinusoidal grating of optimal spatial frequency, length, width and orientation, moving at 4 cycles s^{-1} , presented on a background the same hue and luminance as the rest of the screen. (*b*) Response to the same sinusoidal grating as (*a*) presented on a background shifted towards the blue. The area surrounding the grating remained at constant hue and luminance. The background shift occurred abruptly and coincided with the onset of the grating. The response is bigger than in (*a*). (*c*) Response to the same grating presented on a background shifted away from the blue (towards yellow). The response is bigger than in (*a*). (*d*) Response to the blank screen. (*e*) Response to the background shift towards blue, but with no grating. (*f*) Response to the background shift away from blue, presented with no grating. (*g*) Response to the sinusoidal oscillation of the background between its two extreme colours (towards blue and away from blue) at 4 Hz. (*h*) Response to a sinusoidal grating of optimal spatial frequency, length, width and orientation, which modulates the screen between the two extreme colours of the background and moves at 4 Hz in the same direction as the grating used in figure 6*a–c*.

optimal temporal frequency of either a spatially uniform patch (*g*) or a grating of optimal spatial frequency, orientation and size (*h*). Both when it is spatially modulated at the optimal spatial frequency and when it is spatially unmodulated, the chromatic modulation (in which the colour oscillates sinusoidally between the colours of the bluish and yellowish patches) elicits no discernible response, whereas the response to a luminance modulation of the same temporal and spatial frequency is about 100 spikes per second (figure 6*a*). Thus, the colour-shift




Figure 7. Plots of the extent to which the two backgrounds enhanced the responses of neurons recorded in the striate cortex, area 18 and the LGN. The enhancement is measured as the log of the magnitude of the response in the presence of the background divided by the magnitude of the response in the absence of any background (i.e. the log of the ratio of the two response magnitudes). The enhancement by the blue background is plotted against the enhancement by the yellow background. Cells recorded in the striate cortex are plotted as circles, those recorded in the prestriate cortex are plotted as squares, and those recorded in the LGN are plotted as triangles. For each LGN cell the response magnitude is measured as the magnitude of the Fourier component in the discharge histogram that corresponded to the temporal frequency of the grating. For each cortical cell, the response is measured either as the mean firing rate during the presentation of the grating, or as the magnitude of the corresponding Fourier component, whichever is the larger.

enhances the response to luminance without generating an explicit response of its own.

(c) Summary of colour modulation effects

The non-selective enhancement of neuronal responses illustrated in figure 6 could potentially provide a basis for the neuron to contribute to the discrimination between luminance-only edges and colour edges. In figure 7, we summarize the responses of all the neurons that we have tested with the same set of stimuli by plotting the degree to which the response is enhanced by the blue hue shift, against the degree to which it is enhanced by the yellow hue shift. Neurons whose responses are enhanced by both hue shifts appear in the top right-hand quadrant and neurons whose responses are suppressed by both hue shifts plot in the lower left-hand quadrant. Colour-selective neurons, in which the hue shifts have effects of opposite sign, appear in the other two quadrants. Although the majority of the cortical neurons show no effect, a significant number—three out of the six neurons recorded in area 18 (V2) and three out of the 18 neurons recorded in striate cortex (V1)—show an enhancement of their responses by both hue shifts. One of the neurons recorded in area 18 shows a substantial suppression of its response by both hue shifts.

Also shown in figure 7 are the responses of a larger sample of lateral geniculate neurons in the same experiment. None of the LGN cells shows non-selective enhancement of responsiveness by the hue shift. Most of them lie close to the origin, showing that the response to luminance is not modulated by colour. A small number lie in the lower right quadrant, showing that the two hue shifts have opposite modulatory effects—this is the behaviour to be expected of colour-opponent neurons.

4. ENHANCEMENT OF NEURAL RESPONSES AS A BASIS FOR BEHAVIOURAL DISTINCTIVENESS OF COLOUR

(a) Striate and prestriate cortex

Whatever the function of the perceptual salience of colour, it seems likely that its physiological basis is related to the way that it enhances the responses of neurons early in the cortical visual pathway. The increases in response magnitude caused by the hue shift are clearly detectable in the averaged responses, but modest in size, amounting to no more than about a 30% increase in firing rate in the neurons that we have studied. In the one neuron where we observed a decrease in firing rate, the effect was more dramatic: both hue shifts reduced the response by more than 50%. However, we shall focus on the functional significance of the increases in firing, since they are more common. The two questions to consider are, first, whether these physiological effects are a possible basis for the behavioural results, and second, whether they are the most likely basis for such an effect.

There can be little doubt, in principle, that the physiological response facilitation that we have observed could form the basis for the behavioural discrimination between colour and greyscale images. Although the increases in response magnitude are probably too small to be detected in a single neuron, on a single presentation of a stimulus, the fact that they can be detected in the averaged responses of individual neurons makes it extremely likely that they could be detected in a single trial by averaging across the responses of many neurons. We know that such enhancements are possible: the sensitivity of macaque monkeys in colour discriminations exceeds the sensitivity of the individual neurons that relay colour signals through the monkey LGN (Derrington *et al.* 1984; Derrington 1992). Thus, if we assume that the response enhancements occur across the neuronal population with approximately the frequency that we have observed in our sample, these increases in response magnitude should be easily detectable on a single trial by averaging across neurons.

Of course, only averaging across the population of neurons whose responses are enhanced by colour would not be enough to distinguish colour images from greyscale images. The response enhancements caused by colour could be confused with response enhancements caused by increases in contrast. The results of the luminance transfer test show that the discrimination between colour and greyscale images is independent of the contrast of the greyscale images. To resolve the confusion, the responses would have to be compared with those of similar neurons whose responses were unaffected by colour. Successful discrimination between colour and greyscale images would require

a comparison between the averaged responses of the two populations of neurons.

Consequently, although the neurons that we have recorded could contribute to the discrimination, it is likely that the discrimination itself occurs at a later stage of visual processing, where the signals carried by neurons like those that we have recorded can be pooled. This, of course, raises the possibility that the response enhancements that we record could be a consequence of the colour greyscale discrimination fed back from the later stage, rather than providing the basis for the discrimination that is carried out there.

(b) *Likely involvement of later stages of processing*

Another reason for supposing that the discrimination between colour and greyscale occurs at a relatively late stage in visual processing is the potential role of hue boundaries in image segmentation (Hurlbert 1989). The enhancement of physiological responses associated with hue boundaries that we have observed could be a by-product of the image segmentation process, which is likely to depend on quite late stages of visual processing. One suggestion is that visual processing in cortical area V4 is related, in a general way, to the segmentation of surfaces in the image (Lennie 1998) and that this general segmentation includes aspects of selectivity to colour that are sensitive to differences in hue between the visual stimulus and its surroundings (Zeki 1983*b*). Lesions in V4 appear to disrupt pattern discriminations and complex object-related colour discriminations, but leave simple colour discriminations unaffected (Heywood & Cowey 1987; Heywood *et al.* 1992; Walsh *et al.* 1993). Thus, it is probable that cortical area V4 is central to the physiological and behavioural phenomena described in this paper. It is important to discover both how the responses of neurons in V4 are modulated by hue shifts such as those used in this study, and whether neurons in V1, V2 and V3 still show modulation of their responses by hue shifts when V4 is removed or inactivated.

To answer this question, we need to record from striate and early prestriate cortex in a preparation, with higher cortical areas inactivated or removed. It would also be informative to carry out the behavioural task in a preparation with higher cortical areas removed. In this respect, it is significant that macaque monkeys with inferotemporal cortex lesions show subtle deficits in colour discrimination (Huxlin *et al.* 2000), although macaques have not, to our knowledge, been tested on their ability to generalize a colour/greyscale discrimination.

5. THE USES OF COLOUR VISION

We propose that the utility of colour signals for image segmentation makes colour stimuli distinctive. If so, we should also expect that the availability of colour signals would improve our visual abilities. The evidence on this question is sparse and suggests that colour improves performance but not speed. When monkeys or humans have to classify images according to whether they contain animals or food, colour makes no difference to the speed or accuracy with which they can do the task (Fabre-Thorpe *et al.* 1998; Delorme *et al.* 1999, 2000). On the positive

side, the recognition of natural scenes is improved by the availability of colour signals both during encoding and during recall. The advantage of colour during encoding is interpreted as a sensory effect, and may reflect its usefulness in image segmentation (Gegenfurtner & Rieger 2000).

(a) *Different uses of L-M and short-wave systems*

For convenience, we have concentrated our experiments on the short-wavelength colour-opponent channel by using male marmosets (which have only this channel) for the behavioural work and by targeting our physiological colour stimuli on the short-wavelength channel. In the following paragraphs, we consider whether our conclusions would also apply to the channel carrying signals from the long- and medium-wavelength cones, the L-M channel.

Another reason for concentrating our experiments on the short-wavelength opponent channel is that this is the primordial colour vision system possessed by most mammals that have colour vision (Jacobs 1993). Its function is therefore likely to reflect the primordial role of colour vision. It is clear that the short-wavelength sensitive opponent channel contributes to human memory because, in human X-chromosome-linked dichromats, which lack the L-M (red-green) colour-opponent channel, colour has the same enhancement of memory as it does in normal trichromats (Gegenfurtner *et al.* 1998; Gegenfurtner & Rieger 2000). However, we do not know whether the L-M colour-opponent channel makes such a contribution. This raises the question of whether we would expect the colour images that selectively stimulate the L-M colour-opponent channel of old-world primates to show the same distinctiveness that we have observed in our experiments.

One reason for supposing that the L-M colour channel might not be used in the same way as the short-wavelength channel is that its evolution is more recent (Mollon 1989) and has been linked to the specific colour discriminations associated with fruit eating (Osorio & Vorobyev 1996; Regan *et al.* 1998; Sumner & Mollon 2000). Consequently, we might expect signals in the L-M channel to be used for identifying specific classes of object but not for image segmentation. We have begun experiments with macaque monkeys to test whether the two colour-opponent channels contribute in the same way to the distinctiveness of colour images.

As in most other new-world primate species that have been studied, marmosets have a range of colour-vision phenotypes that can be accounted for by variations in the pigment sensitive to longer wavelengths, which are coded by a single gene on the X-chromosome. There are three such pigments, with peak sensitivities at 545, 556 or 563 nm (Hunt *et al.* 1993). All marmosets share a common short-wavelength-sensitive pigment; males have a single long-wavelength-sensitive pigment and females have two, drawn at random from the three. Thus, all male marmosets and one third of female marmosets are dichromats, with only a single post-receptoral colour-opponent channel (Tovée 1994). Because of the wide variety of colour vision phenotypes and the uncertainty about the extent to which the marmoset visual system is wired to exploit the availability of trichromatic signals in the minority of marmosets that possess them, we did not attempt to investi-

gate any contribution of a possible L–M channel either to physiological responses or to behaviour. However, care was taken to design our stimuli so that the identity of the long-wavelength cone pigment that contributes to the short-wavelength-sensitive colour-opponent channel is irrelevant. Figure 5 shows that the identity of the long-wavelength-sensitive cone makes only a tiny, quantitative difference to the effect of the colour modulations and so variations in cone type would not affect our results.

6. CONCLUSION

Both in their behavioural responses and in the responses of neurons in areas 17 and 18, marmosets distinguish between spatial patterns that contain only variations in luminance and those in which hue also varies. We propose that these distinctions are related to the use of colour signals for image segmentation.

This work was supported by grants from the Wellcome Trust and the BBSRC, a Wellcome Prize studentship to B.S.W. and an MRC studentship to G.R.G. The authors thank Stephen Westland and Sophie Wuerger for their help with colour calibrations, Jackie Francis and Mel Williams for their help with behavioural testing, and Peter Lennie and Anya Hurlbert for the use of their software.

REFERENCES

- Delorme, A., Richard, G. & Fabre-Thorpe, M. 1999 Rapid processing of complex natural scenes: a role for the magnocellular visual pathways? *Neurocomputing* **26**, 663–670.
- Delorme, A., Richard, G. & Fabre-Thorpe, M. 2000 Ultra-rapid categorisation of natural scenes does not rely on colour cues: a study in monkeys and humans. *Vis. Res.* **40**, 2187–2200.
- Derrington, A. M. 1992 Relative sensitivity to colour and luminance gratings in human observers and monkey parvocellular neurones. *J. Physiol.* **452**, 286.
- Derrington, A. M., Krauskopf, J. & Lennie, P. 1984 Chromatic mechanisms in lateral geniculate nucleus of macaque. *J. Physiol.* **357**, 241–265.
- Dias, R., Robbins, T. W. & Roberts, A. C. 1996a Dissociation in prefrontal cortex of affective and attentional shifts. *Nature* **380**, 69–72.
- Dias, R., Robbins, T. W. & Roberts, A. C. 1996b Primate analogue of the Wisconsin Card Sort Test: effects of excitotoxic lesions of the prefrontal cortex in the marmoset. *Behav. Neurosci.* **110**, 872–886.
- Fabre-Thorpe, M., Richard, G. & Thorpe, S. J. 1998 Rapid categorization of natural images by rhesus monkeys. *NeuroReport* **9**, 303–308.
- Felisberti, F. & Derrington, A. M. 2001 Long-range interactions in the lateral geniculate nucleus of the new-world monkey, *Callithrix jacchus*. *Vis. Neurosci.* **18**, 209–218.
- Gegenfurtner, K. R. & Rieger, J. 2000 Sensory and cognitive contributions of color to the recognition of natural scenes. *Curr. Biol.* **10**, 805–808.
- Gegenfurtner, K. R., Wichmann, F. A. & Sharpe, L. T. 1998 The contribution of color to visual memory in X-chromosome-linked dichromats. *Vis. Res.* **38**, 1041–1045.
- Heywood, C. A. & Cowey, A. 1987 On the role of cortical area V4 in the discrimination of hue and pattern in macaque monkeys. *J. Neurosci.* **7**, 2601–2617.
- Heywood, C. A., Gadotti, A. & Cowey, A. 1992 Cortical area-V4 and its role in the perception of color. *J. Neurosci.* **12**, 4056–4065.
- Hunt, D. M., Williams, A. J., Bowmaker, J. K. & Mollon, J. D. 1993 Structure and evolution of the polymorphic photopigment gene of the marmoset. *Vis. Res.* **33**, 147–154.
- Hurlbert, A. 1989 Colour algorithms for image segmentation. In *Models of brain function* (ed. R. M. J. Cotterill), pp. 89–113. Cambridge University Press.
- Huxlin, K. R., Saunders, R. C., Marchionini, D., Pham, H. A. & Merigan, W. H. 2000 Perceptual deficits after lesions of inferotemporal cortex in macaques. *Cerebr. Cortex* **10**, 671–683.
- Jacobs, G. H. 1993 The distribution and nature of colour vision among the mammals. *Biol. Rev. Camb. Phil. Soc.* **68**, 413–471.
- Lennie, P. 1998 Single units and visual cortical organization. *Perception* **27**, 889–935.
- Marr, D. 1976 The early processing of visual information. *Phil. Trans. R. Soc. Lond. B* **275**, 483–524.
- Marr, D. 1982 *Vision: a computational investigation into the human representation and processing of visual information*. New York: W. H. Freeman & Co.
- Marr, D. & Hildreth, E. 1980 Theory of edge detection. *Proc. R. Soc. Lond. B* **207**, 187–217.
- Mollon, J. D. 1989 “Tho’ she kneel’d in that place where they grew. . .” The uses and origins of primate colour vision. *J. Exp. Biol.* **146**, 21–38.
- Osorio, D. & Vorobyev, M. 1996 Colour vision as an adaptation to frugivory in primates. *Proc. R. Soc. Lond. B* **263**, 593–599.
- Regan, B. C., Julliot, C., Simmen, B., Vienot, F., Charles-Dominique, P. & Mollon, J. D. 1998 Frugivory and colour vision in *Alouatta seniculus*, a trichromatic platyrrhine monkey. *Vis. Res.* **38**, 3321–3327.
- Stavenga, D. G., Smits, R. P. & Hoenders, B. J. 1993 Simple exponential functions describing the absorbency bands of visual pigment spectra. *Vis. Res.* **33**, 1011–1017.
- Sumner, P. & Mollon, J. D. 2000 Catarrhine photopigments are optimized for detecting targets against a foliage background. *J. Exp. Biol.* **203**, 1963–1986.
- Tovée, M. J. 1994 The molecular genetics and evolution of primate colour vision. *Trends Neurosci.* **17**, 30–37.
- Tovée, M. J., Bowmaker, J. K. & Mollon, J. D. 1992 The relationship between cone pigments and behavioural sensitivity in a new world monkey (*Callithrix jacchus jacchus*). *Vis. Res.* **32**, 867–878.
- Travis, D. S., Bowmaker, J. K. & Mollon, J. D. 1988 Polymorphism of visual pigments in a callitrichid monkey. *Vis. Res.* **28**, 481–490.
- Walsh, V., Carden, D., Butler, S. R. & Kulikowski, J. J. 1993 The effects of V4 lesions on the visual abilities of macaques—hue discrimination and color constancy. *Behav. Brain Res.* **53**, 51–62.
- Wiesel, T. N. & Hubel, D. H. 1966 Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *J. Neurophysiol.* **29**, 1115–1156.
- Zeki, S. 1983a Color coding in the cerebral-cortex—the reaction of cells in monkey visual-cortex to wavelengths and colors. *Neuroscience* **9**, 741–765.
- Zeki, S. 1983b Color coding in the cerebral-cortex—the responses of wavelength-selective and color-coded cells in monkey visual-cortex to changes in wavelength composition. *Neuroscience* **9**, 767–781.

GLOSSARY

LGN: lateral geniculate nucleus