

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank) 2. REPORT DATE 31 Dec 1990 3. REPORT TYPE AND DATES COVERED Final report 1 Nov 89 - 31 Oct 90

4. TITLE AND SUBTITLE Neurophysiological Implementation of a Scheme for Visual Selective Attention 5. FUNDING NUMBERS G AFOSR-90-0043 PE-61102F PR-2313 TA-A4

6. AUTHOR(S) Dr. John Duncan

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Medical Research Council Applied Psychology Unit 15 Chaucer Road Cambridge CB2 2EF United Kingdom 8. PERFORMING ORGANIZATION REPORT NUMBER AFOSR-TR-91-0138

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research Bolling AFB DC 20332-6448 10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES Prepared in collaboration with E. Miller, L. Chelazzi, J. Hart and R. Desimone, Laboratory of Neuropsychology, NIMH, Building 9, Room IN107, Bethesda, MD 20892

12a. DISTRIBUTION/AVAILABILITY STATEMENT Distribution unlimited 12b. DISTRIBUTION CODE DTIC ELECTE MAR 11 1991

13. ABSTRACT (Maximum 200 words) The method of single cell recording in the behaving monkey was used to study mechanisms of selective attention in extrastriate cortex. The animal indicated whether any item in a display of one or more colour patches matched a previously presented target. Recordings were obtained from 186 single cells in cortical area V4. In the match task, attentional selection of the target patch must be controlled by an advance description of the target colour. No evidence for such enduring target descriptions was found in V4: Few if any units showed sustained, target-specific activity. While there was clear evidence for attentional modulation of the visual response, this modulation was quite different from that previously reported for a location selection task. Units did not always respond more strongly when attention was directed to a patch of their preferred colour; instead the distribution of attentional preferences was bimodal, with roughly half the cells giving stronger responses when their preferred colour was ignored. Such effects, furthermore, were dependent only on preferences for colour (the relevant dimension); location preferences seemed immaterial. The findings are inconsistent with hypotheses for attentional selection based on either input gating or selective priming. Alternative possibilities are discussed.

14. SUBJECT TERMS Vision, attention, visual neurophysiology, single unit recording 15. NUMBER OF PAGES 52 16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT Unclassified 18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified 19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified 20. LIMITATION OF ABSTRACT Unlimited

AL 7074 040

2

Background

[illegible]

QUALITY
INSPECTED
3

91 3 06 206

attentional modulations arise, and how they produce a state in which only sensory input from relevant or attended objects gains control of behaviour.

The attentional template

In human vision, many different stimulus attributes can be used to guide attention to relevant objects. For example, when choosing which items to report from a brief display, people can select rather efficiently on the basis of spatial location (e.g., attending to letters in a particular row), size, motion, brightness, shape category (e.g., alphanumeric class), etc (Duncan, 1983; Merikle, 1980; von Wright, 1968). Such flexibility is what we might expect of a useful attentional system, since tasks vary widely in their specification of what sort of visual information is "relevant".

The implication is that selection must be controlled by some sort of variable advance description of the sort of information needed. We call such an advance description the attentional template (Duncan & Humphreys, 1989). The sensory description of objects present in the visual field must be compared against such a template, allowing selection of those objects that match. In the earlier example, the template might specify simply "red". When the display is then presented, objects whose colour matches this template are selected, while mismatching objects are gated out.

How might such templates be implemented neurophysiologically? One possibility is suggested by the specialisation of different extrastriate regions for the analysis of different visual attributes. For example, a pathway leading ventrally from striate cortex through areas V2 and V4 into inferotemporal cortex seems specialised for the analysis of shape and colour, hence ultimately object recognition (Desimone & Ungerleider, 1989). In contrast, a pathway leading dorsally into the parietal lobe is more concerned with analysis and representation of spatial locations. One hypothesis might

be that different attentional templates are represented in different extrastriate regions, depending on their content. For example, colour templates might be represented in one or more areas of the ventral visual stream, while spatial templates might be represented in the dorsal stream.

Results in support of this hypothesis were reported by Haenny, Maunsell, & Schiller (1988). In their experiment, the animal was trained in a cross-modal match to sample task. At the start of each trial, a sample or target orientation was specified by touch, and the animal's task was then to watch a sequence of visually-presented gratings at different orientations, releasing a lever when a match to the target eventually occurred. Recording from single cells in V4, Haenny et al. (1988) found the expected modulation of neural activity by current visual input. For example, a given cell might fire most strongly when vertical lines were presented in the receptive field. Just as striking, however, was modulation by the target orientation for a given trial. Irrespective of the current visual input, for example, a neuron might fire most strongly for trials when the required target was vertical; and in some cases, such discharges began at the time the target was first specified and continued until the trial was complete. Such results strongly suggest coding of target descriptions, used for match against later visual inputs, even as early in the visual system as V4. Similarly, it has been reported that some neurons in area LIP of the parietal cortex can hold a short-term representation of the target location for forthcoming behaviour (Gnadt & Andersen, 1988).

There are however plausible alternatives to the idea that attentional templates are coded in extrastriate regions. For example, activity coding the target location for a forthcoming eye movement is also found around the principal sulcus in the frontal lobe (Funahashi, Bruce, & Goldman-Rakic, 1989). Different regions of extrastriate cortex are reciprocally connected with

corresponding regions of the frontal lobe, so that frontal sites too are plausible candidates for template activity.

In the present work, we investigated the extrastriate hypothesis by looking for activity suggesting colour templates in area V4, whose partial specialisation for colour analysis is well known (Zeki, 1980).

Modulation of the visual response

In several extrastriate regions, as we have noted, visual responses are altered by the animal's attentional state. In area 7a of the parietal lobe, for example, visual responses are enhanced when the animal must pay attention to a stimulus presented in the receptive field, monitoring it to detect a change (Bushnell et al., 1981). In an experiment by Moran & Desimone (1985), pairs of stimuli (coloured bars) were presented within the receptive field of single neurons of both V4 and IT. For each neuron, stimuli were chosen such that one stimulus (the preferred stimulus) would excite the cell while the other would not. Using a form of spatial cueing, the animal's attention was shifted between blocks of trials from one stimulus to the other. The results showed responses determined almost entirely by the attended stimulus, cells responding well when their preferred stimulus was attended but much less well when it was ignored.

The details of this effect are also important (see also Wise & Desimone, 1988). Cells always responded well when their preferred stimulus was presented alone in the receptive field, whether or not it was behaviourally relevant. Thus the results with stimulus pairs suggest a specific inhibition of response to the preferred stimulus when attention was paid to another, nearby object. Furthermore, this inhibition did not begin with the onset of the visual response, which in V4 is typically 40-60 msec after onset of the stimulus. Even when the cell's preferred stimulus was the one that the

animal ignored, the visual response developed normally up to around 90 msec poststimulus, then was abruptly cut off. It seems sensible to suggest that the lag between onset of the visual response and onset of attentional inhibition reflected the time needed for comparison of the visual input with some form of spatial template specifying (for this task) which bar was relevant.

In this report we consider three different models for how attentional modulations of the visual response might be produced. In each case the guiding principle is as follows. Functionally, the goal of visual attention is presumably to ensure that the relevant or attended object controls behaviour. For an extrastriate area like V4, what does this imply? Neurons have sensory preferences along numerous different visual dimensions (location, colour, orientation etc). The obvious suggestion is that, once the attentional state has developed, all neurons whose sensory preferences match the properties of the attended object should remain active, while all those with other sensory preferences should be silent. Then, irrespective of how the object had originally been selected (e.g., based on its location, colour etc), any subsequent system receiving the area's output would be sent a message reflecting all and only the properties of the selected object. All three of the models that we consider share this same general characteristic.

Model 1 may be called input gating. As pointed out by Moran and Desimone (1985), one hypothesis is that, when an attentional state develops, the receptive fields of cells in V4 or IT effectively shrink so as to encompass only the attended region. If the preferred stimulus - the stimulus actually capable of exciting the cell - lies within this region then the response remains, but otherwise it is cut off. According to this model, selection works on the

inputs to V4, gating out all those from objects lying outside the attended region (see also e.g., Feldman, 1985).

Model 2 may be called biased competition. Suppose that attentional templates are implemented by priming or preactivating cells with the corresponding sensory preference. For example, when the animal searches for a red target, cells preferring red on the colour dimension are preactivated. The results of Haenny et al. (1988) indeed suggest an approximation to this pattern: While in terms of orientation a V4 cell's target and visual preferences could be different, there was a strong bias towards their being the same. For example, if a cell fired most strongly on trials when the target orientation was vertical, then the most likely visual preference was also for vertical. Suppose in addition that cell populations with different sensory preferences (e.g., for red and green) are mutually inhibitory. Then priming a particular population would give them a competitive advantage in this mutual inhibition; and when we recall that cells typically have multidimensional preferences (e.g., for a red vertical at a given location), it can be seen that the cells most active after such a competitive process will be those coding the particular conjunction of properties possessed by the target object (cf. Phaf, van der Heijden, & Hudson, 1990).

The third model may be called feedback modulation. Without question attentional selection involves both cortical and subcortical structures outside the prestriate region. Quite possibly, the attentional effects seen for example in V4 are merely the result of a selection process taking place elsewhere. In this case we might see attentional modulation as an independent, perhaps late input to V4, returning from other brain structures and combining with the sensory input to produce the overall response.

By investigating the details of attentional modulation in V4 in a colour selection task, we hoped to cast light on these three alternative hypotheses.

Outline of the method

In human vision, one of the most effective attentional selection cues is colour. For example, a person trying to report only red letters from a display is scarcely influenced by the simultaneous presence of black letters (Bundesen, Shibuya, & Larsen, 1985). Similarly, if a person is asked to search a display of colour patches to decide whether a particular colour (e.g., a red patch) is present, then providing colours are not very similar, search time is independent of the number of (wrong-coloured) nontarget patches (Duncan, 1989). This is a task giving a very strong attentional "popout": Irrespective of the number of nontargets, attention is drawn directly to the target. Accordingly we decided to investigate an analogous colour search task in the monkey, looking at the activity of individual V4 neurons for evidence of both template activity and modulations of the visual response.

In our task, each trial began with a foveal colour patch (the sample) indicating the target colour for this trial. Around a second later, there was a test display containing either one or two peripheral patches, presented in the region of the recorded cell's receptive field. If there was a patch of the target colour, the animal released a lever for reward. Otherwise he held on for reward later in the trial. Unbroken central fixation was required throughout.

Neural responses were analysed during three phases of each trial. First consider responses to the sample. Given our recording site, for most cells the (foveal) sample patches were presented well away from the centre of the receptive field. Accordingly we expected little response in a control task, involving central colour patches that were not samples for a subsequent peripheral test. Following the results of Haenny et al. (1988), however, it

seemed possible that we would see strong, colour-selective responses at the time of the foveal patch in the main task. Such findings might be one index of establishing colour templates for a search task in V4.

A second feature of Haenny et al.'s (1988) results was that responses selective for target orientation were sometimes sustained throughout a trial. Secondly, therefore, we analysed neural activity in the interval between sample and test, asking whether there was sustained activity, selective for sample colour, even after any simple visual response to the sample was over.

Thirdly we analysed visual responses to the test displays themselves. Our main interest centred on the case in which displays contained two patches, one in a cell's preferred colour, the other in a different colour. Following the results of Moran and Desimone (1985), we might predict larger responses to such a display when attention was focussed on the preferred colour patch (i.e., this patch matched the sample) than when attention was focussed on the other patch. Their results, however, were obtained with a spatial selection cue, and there is reason to suppose that other sorts of cue could behave rather differently. Behaviourally, spatial selection is especially effective (von Wright, 1968), and its influence on the visual evoked potential is both earlier and different in form from the effects of selection based on other visual attributes (Hillyard, Munte, & Neville, 1985). Thus we wanted to ask whether results like those of Moran and Desimone (1985) would be obtained in a colour selection task, and if so, to investigate the three alternative models for attentional modulation.

METHOD

General procedure

The main data described in this report were obtained in recordings from an area provisionally identified as the ventral division of V4 (upper

field representation) in the right hemisphere of one male macaque. Supporting data are currently being collected from dorsal V4 in a second animal. Responses of isolated neurons were recorded during performance of various colour match tasks.

The experiment was controlled by a 386 PC computer running a software package developed at NIMH for electrophysiological studies. The monkey was restrained in a primate chair, with head fixed by means of a bolt cemented to the skull and eye position continuously monitored by magnetic search coil. Displays were presented on a colour monitor, and responses signalled on a hand held bar.

Extracellular recordings were made with paralene coated microelectrodes. For each penetration, a guard tube containing the electrode was lowered through the dura, and the electrode was then advanced from within it. Ventral V4 was provisionally identified from MRI scans, and the recording well was mounted stereotaxically so that the electrode, advanced vertically down through the cortex, would reach the intended region. Recordings were made in the last bank of cells encountered before exiting the base of the brain. The topology of the visual representation in this region supports the view that it is ventral V4. As penetrations move anteriorly, the preponderance of receptive fields moves from the vertical towards the horizontal meridian, as expected for V4 but not the adjacent areas V3 and TEO. Since the animal has not yet been killed, however, there is still no final histological confirmation of the recording site.

Single neuron activity was extracted from the pooled multiple neuron signal by an on-line spike sorter using a template matching algorithm (Signal Processing Systems). In many cases it was possible to identify two separate neurons from the same site. Pulses from the spike sorter were recorded by

the computer along with codes indicating trial events. When a neuron was isolated, its spatial receptive field and colour preferences were roughly mapped, and responses were then recorded during one or more colour match tasks. Neurons were only accepted that appeared to give a positive response to task stimuli. Recordings from a given neuron usually lasted about an hour. In each session, fresh neurons were isolated and recorded until the animal stopped working, generally after around three hours.

Three colour task

For most cells the main recordings were done in a three colour task. In nearly all cases the three colours used were red, green and blue presented on a dark grey background, though occasional cells that seemed strongly to prefer a different colour were given instead orange, bluish-green and purple. Colours for any given cell were closely matched in luminance; across cells luminance varied in the range $10\text{-}20\text{ cd/m}^2$.

Each trial began with the monkey resting his hand on the response bar. A small white fixation point appeared in the centre of the screen, and 300 msec after adequate fixation had been achieved (plus or minus 0.5°), a 1° -diameter circular colour patch (the sample) was presented at the fovea. It remained for 200 msec, and was followed after a randomly varying interval of 700-1200 msec by a peripheral test display. This display contained either one or two colour patches, also 1° in diameter, and also lasting for 200 msec. If the test display contained a patch matching the sample colour, then the animal was supposed to release the response bar for immediate auditory feedback and juice reward. Otherwise, there was a further random interval of 700-1000 msec, followed by a single peripheral patch which was always a match to the sample and which always indicated that the bar should be released for reward. Unbroken central fixation (criterion as above) was

required throughout the trial. Loss of fixation, premature bar release, or late release (maximum permitted latency 550 msec) all caused the trial to be immediately aborted without reward. The minimum interval between trials was approximately 600 msec.

For any given neuron, two peripheral locations were chosen for patches in the test display. One - which we shall call the x location - was chosen to be near the centre of the cell's excitatory receptive field, as very roughly determined in initial observations, and the second (y location) was at the same eccentricity but separated by 45° (or in a few cases either 25° or 90°) in polar coordinates from the first. Because of our recording site, x locations were always in the upper left quadrant of the visual field, with eccentricity between 2.2° and 9.9° visual angle.

The experimental design is shown in Table 1, which assumes the usual case in which the three colours used were red, green and blue. When the test display contained only a single patch, it could be in any of the three possible colours and either the x or y location. Combination of these six possible test displays with three possible samples gives eighteen single patch conditions. When the test display contained two patches, they always differed in colour, with one in the x and one in the y location. Again this gives six possible test displays preceded by three possible samples, for a total of eighteen two patch conditions.

Insert Table 1 about here

For the majority of cells, all 36 conditions occurred in random order, and equally often except that the two patch mismatch conditions (31, 32, 27, 28, 23, 24) each occurred twice as often as the remainder (so that only half of

the two patch trials would be matches). For a few cells recorded at the start of the experiment, however, trials were run in blocks of 24 with the same sample colour. Recordings with a given cell were continued until approximately ten correct trials had been obtained for each condition.

For one patch trials the animal typically performed at over 85% correct. For two patch trials, however, accuracy was generally no better than 70 to 80%.

Foveal control task

In some cases after recordings had been completed in the three colour task, data were also collected for a foveal control task. It was exactly like the main task, except that the test display always contained only a single patch, presented like the sample at the fovea. Combination of three possible sample colours with three possible test colours gave a total of nine conditions, which occurred equally often in random order.

Six colour task

In a subsidiary experiment, recordings were made in a six rather than a three colour task. The six colours - red, yellow, green, acqua, blue, purple - were chosen to be about equally spaced around the conventional colour circle, given the constraints of the colour monitor.

Colours were divided into two groups of three - red, green, blue and yellow, acqua, purple - and for each group the design was a replica of that shown in Table 1. Thus each single patch occurred equally often, but only certain two-patch combinations were possible (e.g., red appeared with green and with blue, but not with the other colours). The resulting 72 possible conditions again occurred in random order, and again recordings were continued to a criterion of about 10 correct trials per condition.

Preparation of the animal

Training and preparation of each animal took at least 12 months before recordings began. The animal was trained first simply to release the bar for juice reward; then to release when a stimulus appeared on the screen; then to ignore a first stimulus (progressively increased in duration) and wait for a second, matching stimulus; then to release only when the second stimulus matched the first (cf. the foveal control task above). At this point eye coil, head bolt and recording chamber were fitted surgically under aseptic conditions. The monkey was administered intramuscular ketamine and atropine, and deeply anaesthetised with intravenous pentobarbitol. Following surgery, the animal was given Tylenol to reduce pain and antibiotics as a prophylactic against infection. Training resumed after recovery. The test display was moved off the fovea, and the criterion for maintaining central fixation was gradually tightened. Finally the two patch test display was introduced, and training continued until performance was better than 80% correct. As noted above, however, performance was usually somewhat worse in recording sessions.

RESULTS

Three colour task

Recordings were obtained from 186 cells, 140 of which showed a clear positive response to peripheral displays. All analyses took account only of data from trials on which the monkey's response was correct.

Sample responses

While large responses to the foveal sample were rare, small responses were often seen. In 106 cells these could be compared with corresponding responses to foveal sample and test displays in the foveal control task. These comparisons revealed no obvious trend and very few large differences. Since

many artefacts could influence the comparison of responses separated by an average of perhaps half an hour, these data were analysed no further. The conservative conclusion is that we obtained no striking evidence for responses to the foveal sample that were any different from the simple visual response seen when foveal patches were not samples for a subsequent peripheral test.

Delay activity

On all trials there was an interval of at least 700 msec between offset of the sample and onset of the test display. To avoid possible off-discharges to the sample the first 160 msec of this interval was discarded. The remaining 540 msec was divided into three successive subintervals of 180 msec each, and for each cell, response rates (expressed in impulses/sec) were subjected to analysis of variance (ANOVA) with sample colour (three levels) and subinterval (three levels) as factors. The analysis was restricted to the 174 cells recorded with sample colour completely randomised across trials.

Selective information about target colour for the current trial could be carried by either a main effect of sample or an interaction between sample and subinterval. The main effect was significant ($p < .05$) for only 13/174 or 7.5% of cells, while the interaction was significant for only 16/174 or 9.2%. Bearing in mind that each effect would be measured as statistically significant in 5% of cells by chance alone, it seems safe to conclude that in our task few if any V4 cells have measurable delay activity that could be related to the coding of an attentional template. It should also be remembered that our analysis was comparatively powerful, since it was typically based on over a hundred trials per sample colour.

Figure 1 shows sample and delay activity for a cell with significant effects of both sample colour ($p < .001$) and the interaction with subinterval (p

< .001). Time 0 indicates sample onset, 200 indicates offset, and 900 indicates the earliest possible test onset. (Because of the random interval separating sample and test, responses to the test itself - which in this cell were very weak - are smeared over the range beyond 900 msec.) A positive response around sample offset is followed by prolonged inhibition for a red sample, little effect for a green sample, and inhibition followed by prolonged facilitation for the blue sample. Even for this cell, however, the firing rate has returned to the same value for all samples before the test display actually arrives.

Insert Figure 1 about here

Test responses

Form of the response Figure 2(a) shows the response of a typical cell to the test display (onset at time 0). The visual response has a latency of around 50 msec, and takes the form of a sharp initial peak followed by a sustained discharge until well after test offset. While responses of this general form were very common, others also often occurred. Figure 2(b) shows a cell whose response climbs throughout the display interval, while in 2(c), a sharp initial peak is followed by a second peak before test offset (and a third representing an off discharge).

Insert Figure 2 about here

Except where noted, the following conventions were adopted for subsequent analyses. For the assessment of sensory preferences we took mean firing rates (impulses/sec) in the interval between 40 and 220 msec following display onset. For the assessment of attentional modulations,

however, we took mean firing rates between 130 and 220 msec, since the expectation is that such effects would develop only some time after the visual response has begun. Analyses were restricted to the 140 cells showing a positive visual response.

Attentional modulation: Colour To recapitulate, the major test of attentional modulation is made on response to a test display containing two patches, one of which (the preferred stimulus) excites the cell more strongly than the other. Following the human data, we assume that attention tends to focus on the target or match stimulus in the test display. Accordingly, we compare responses to the same display of two patches, for trials on which either the preferred or the alternative patch is the target.

Our first analysis concerned preferences for colour. The first step was to assess such preferences in responses to single patches. Mean responses to each colour were obtained by averaging together the appropriate conditions, e.g., for red, conditions 1, 2, 7, 8, 13, 14 (see Table 1), and various criteria were used to assess colour preference. For this report the criterion adopted is that response to the best colour (the b colour) should be at least 1.3 times response to the worst colour (the w colour). (We call the remaining middle colour the m colour.) This criterion corresponds closely to taking all cells for which the difference between b and w colours was significant at $p < .10$ in ANOVA; a more conservative criterion gives much the same results with fewer cells. Using the present criterion, we obtained a population of 79 colour selective cells.

For each selective cell we then calculated an attentional modulation index, based on all two patch match displays that contained the b colour. The index is defined as

$$\frac{\text{mean response given b sample}}{\text{mean response given other sample}}$$

For example, if the b colour was red, then the index would be calculated as the (unweighted) mean response for conditions 19, 20, 21, 22 divided by the mean for 25, 26, 33, 34 (see Table 1). Expressing the index as a log to base 2 gives us a scale in which 0 indicates no modulation, 1 indicates a response that is doubled when the b colour is the target, -1 a response that is halved.

From the findings of Moran and Desimone (1985) we predicted a distribution of strong positive modulation indices. The results, shown in Figure 3, were quite different. Across cells, the distribution has no positive skew. Instead it is symmetrical around zero, and clearly bimodal. For about half the cells (positive cells; $N = 37$) the response is enhanced when attention is focussed on the b patch, but for the remainder (negative cells; $N = 42$) the response is reduced in this case.

Insert Figure 3 about here

To assess the significance of these results, we asked whether the modulation index tends consistently to be either positive or negative for a given cell. For each cell we calculated the index twice, once for displays containing the b and the w colour, and again for displays containing b and m. The resulting two indices had the same sign in 53/79 (67%) of cells, chi squared = 9.16, $p < .005$.

A more complete summary of responses to two patch displays is presented in Table 2. Mean response rates are presented for all two patch

displays (averaging across different spatial arrangements of the same colours), and for all samples. Which colours were b, m and w was of course separately determined for each cell, and cells have been split into positive and negative sets based on the overall modulation index. Data again are mean firing rates for the interval 130-220 msec after test onset.

Insert Table 2 about here

Several points are worth noting. First, in both positive and negative cells, modulation was roughly similar for b+w and b+m displays. Second, responses on mismatch trials tended to be intermediate between the two types of match trials. One can imagine many possible reasons for this - for example, attention could be randomly divided between the two patches when neither matches the sample - and it will not be discussed further. Third, neither for positive nor for negative cells was there any suggestion of an effect of sample in the m+w displays. Response rates in this case were the same whatever the preceding sample. The attentional effect seems sharply localised to displays containing the b colour.

A further analysis expands this conclusion. To summarise results so far: When attention is directed to the b colour, response is enhanced for some cells, but inhibited for others. Might there also be other, perhaps weaker, effects of attending to other colours, which could also be positive or negative? More generally, might there be some underlying broad function of enhancement and inhibition effects produced by attending to different regions of colour space?

This model is easily tested and refuted. Across cells it predicts that, if the modulation ratio is more positive in the b+w than in the b+m display,

then it should also tend to be positive (preference to attend m) in the m+w display; and vice versa. The prediction follows since, for example, a positive index for m+w would imply more facilitation (or less inhibition) for attending to m than to w. Across the population of 79 cells there was in fact no tendency of this sort. We conclude that attentional modulation is rather local in colour space, arising (given the present selection of colours) only close to the region of the b colour.

Before leaving the two patch data, we may ask whether the effect we have seen is indeed confined to the later part of the visual response. In Figure 4, modulation indices are plotted for the interval 70-130 msec post onset. Data are shown separately for positive and negative cells (as defined by the 130-220 msec index). The results suggest no more than a hint of an attentional effect in the earlier interval; distributions for both positive and negative cells are unimodal around zero. Any slight skew in these distributions suggests that modulation perhaps begins a little before 130 msec; the main effect, however, is later.

We may also ask whether responses to single patch displays were influenced by the preceding sample. Mean response rates are shown in Table 3, again for the interval 130-220 msec post onset and with positive and negative cells (as defined by two patch modulation) separated. In no case is there any real suggestion that responses were influenced by the sample. For example, in neither positive nor negative cells was there any preference for match vs. mismatch trials, or for the b vs. m or w sample. The largest effect of sample occurred for responses to the b test in negative cells, but even here, the response given a b sample was greater than the mean given m and w samples for only 23/42 cells.

Insert Table 3 about here

Further analyses confirmed this conclusion. The first of these looked at data from the full response interval, 40-220 msec post onset. For all 140 visually responsive cells, responses to single patches were examined in a three-way ANOVA, with factors test colour (3 levels), sample colour (3 levels), and location (2 levels). The main effect of sample colour was significant ($p < .05$) in 6/140 or 4.3% of cells, while the interaction with test colour was significant in 7/140 or 5%. These values do not exceed chance expectations. The final analysis was based on the following reasoning. Two patch data suggest that colour selective cells receive a late input, positive for some and negative for others, when attention is paid to the b colour. Does response to a single patch of the b colour also reflect such an input? If so, then in positive cells the late response to the b colour should be enhanced in comparison to responses to m and w colours, while in negative cells late response to the b colour should be inhibited. To put this another way, for positive cells colour tuning should become sharper late in the response, while for negative cells it should become flatter. For all colour selective cells, the strength of colour preference was measured separately for the intervals 40-100 and 130-220 msec. No evidence for the prediction was obtained. We conclude that the modulation we have observed is restricted not only to the case of paying attention to the b colour, but also to displays containing two different patches, i.e., potential attentional competitors.

In Figure 5, average responses to a single patch of the b colour are compared with mean responses to b+w and b+m match displays, separated according to whether the b or the alternative colour is the match. Histograms

were first plotted separately for each cell, and then averaged across cells. The results are striking. For both positive and negative cells, the early part of the response is roughly equal in all three conditions. Note the implication that initial response to a two patch display is roughly comparable to response to the single best patch that it contains. Beginning around 100 msec, however, inhibition begins to develop to two patch displays, which for positive cells is stronger when the b colour is unattended, while for negative cells it is stronger when the b colour is attended. Inhibition dissipates by the time of stimulus offset, when again responses in the three conditions converge.

A final analysis concerns the relationship between strength of attentional modulation and strength of colour preference. Do cells with stronger sensory preferences also show stronger attentional modulations? The answer was clearly negative. Across the 37 positive cells, the rank correlation (ρ) between sensory preference for b over m or w and modulation index was $-.14$; across 42 negative cells it was $.02$; taking both groups together it was $-.09$. (For this analysis only, each cell's measure of sensory preference was based on the same interval as the modulation index, 130-220 msec post onset.)

Attentional modulation: Location In our sample, most of the 140 responsive cells gave some positive response to single patches in both x and y locations. In other words, spatial receptive fields spanned at least 45° in polar coordinates, and presumably usually rather more since the x location was chosen to be roughly at the centre of the receptive field with the y location 45° to one side or the other. Nevertheless, many cells showed a clear location preference, and for these we can ask the same question as for colour: Is response to a two patch array dependent on whether attention is focussed on a patch in the more or less preferred location?

To test this we confined attention to those cells recorded with x and y locations separated by 45° , and selected 61 cells with a significant location preference in ANOVA ($p < .05$). For two patch match trials we defined a new attentional modulation index as

$$\frac{\text{mean response given attention to better location}}{\text{mean response given attention to worse location}}$$

For example, if the preferred location was x, then the index was given by mean response in conditions 19, 26, 21, 34, 29, 36 divided by mean response in conditions 25, 20, 33, 22, 35, 30 (see Table 1). Again the index was based on the interval 130-220 msec post onset of the test display.

Results are shown in Figure 6. They are clearcut: Response to two patch displays is independent of whether the matching (attended) patch occurs in the cell's preferred location. There is no hint of the bimodal distribution seen in colour data.

Insert Figure 6 about here

Results were the same when positive and negative colour cells (as earlier defined), as well as cells with location but not colour preference, were analysed separately. The negative results were not due to location preferences being weaker than colour preferences; for the cells shown in Figure 6, the median location preference (response to better location divided by response to worse location, based on single patch data) was 1.52, as compared to a median colour preference of 1.59 for the cells in Figure 3. Finally, again, there was no correlation across cells between modulation index for location and strength of location preference, $\rho = .15$.

Six colour task

To summarise: When attention is focussed on a stimulus of a cell's preferred colour, the simple sensory response is combined with an additional, late attentional input. For some cells this input is positive, while for others it is negative. The six colour study was run to test two alternative hypotheses concerning the sign of this late attentional input.

Insert Figure 7 about here

These alternatives are illustrated in Figure 7. At the top is shown a hypothetical sensory response, with maximal response to some preferred colour and progressively reduced response to more and more distant or dissimilar colours. For example, the cell might respond maximally to green, somewhat less to yellow and acqua, and so on. Beneath are shown possible attentional inputs, which we suppose are added to the sensory response between 100 and 200 msec after stimulus onset. According to Model 1, there are two alternative attentional inputs, either positive or negative. Each is sharply localised around the cell's preferred colour, and arises only when attention is focussed on a patch of that colour. According to this model, a cell is characteristically either "positive" or "negative" for our task and stimuli. In Model 2, by contrast, all cells behave in exactly the same way, and whether a given cell appears "positive" or "negative" depends essentially on the accidental choice of stimulus colours. The idea here is that the same late attentional input is received by all cells; this input is positive when attention is focussed on a patch closely matching the cell's sensory preference, negative for colours outside but still close to the preferred colour zone, and neutral for more distant colours. For example, if a cell's preferred colour were green, the

attentional input could be positive for green, negative for yellow and acqua, and neutral for all other colours. The result would be an attentional tuning curve with a characteristic "Mexican hat" shape plotted in colour space. For the three colour experiment, a cell would appear "positive" when its favourite among the three colours actually used happened to fall within its central, positive zone, while it would appear "negative" when this favourite actually lay in the surrounding negative zone.

The two models can evidently be distinguished by an experiment using more than three colours. According to Model 2, the proportion of "positive cells" should be appreciably higher when six colours are used, since the chance of finding the cell's central, positive zone should be substantially increased. Indeed, plausible estimates of how wide "positive" and "negative" zones would have to be, given the observed positive and negative modulations in the three colour study, suggested that essentially all cells should be positive with six colours.

Sensory preferences

Of 54 cells recorded in the six colour task, 19 showed a clear positive response to the test display and a significant colour preference ($p < .05$ by ANOVA on all single patch responses). Analyses were conducted on just the data from these 19 selective cells.

Colour tuning curves, based on responses to single patches in the interval 40-220 msec post onset, are illustrated in Figure 8. These curves reflect sensory preferences; they are based on average responses to test patches of different colours, irrespective of the colour of the sample. For the purposes of the figure, colours are always plotted so that movement clockwise round the conventional colour circle (red-yellow-green-acqua-blue-purple-red) goes from left to right along the x axis, and the adjacent pair of

colours giving the best average response is plotted in the centre. The results may be summarised as follows. For many cells there were two adjacent colours (e.g., blue and purple) giving a good response. Response was least to colours on either side of this central excitatory zone (for the example, acqua and red), with a slight upturn for even more distant colours (green and yellow). Figure 8a shows a typical example. In some cases the upturn was so strong that a cell's second favourite colour was opposite to the favourite (e.g., favourite red, second acqua); an example is shown in Figure 8b. The average tuning curve for all 19 cells is shown in Figure 8c, obtained by plotting data for individual cells as in the upper panels and then taking means. The upturn for colours most distant from the favourite is clearly seen, and in fact was observed for 16/19 cells. For the sensory response, therefore, colour tuning curves do have a standard "Mexican hat" shape.

Insert Figure 8 about here

Analysis of attentional modulations, however, rules out the hypothesis that for attentional inputs too there is a central positive zone and surrounding negative zones. Again analyses were based on two patch displays containing the cell's most preferred colour, and conducted as before on responses in the interval 130-220 msec post onset. Of the 19 cells, 9 were positive and 10 were negative. The direction of modulation was the same in b+m and b+w displays for 11/18 or 61% of cells (one cell missing since modulation index was precisely 1.0 for one display). These data are similar to those from the three colour task, and certainly give no hint of a shift towards a preponderance of positive cells. Model 2 in Figure 7 is ruled out.

Post-stimulus inhibition

Though not directly related to our main concerns, one other aspect of the data deserves mention. This is a marked and prolonged inhibition that seems to develop in V4 around 200-400 msec post stimulus offset.

One clue of such an effect is provided by our analysis of responses in the interval between sample and test. Recall that for this analysis we discarded the first 160 msec post sample offset, then examined firing rates in three subsequent subintervals of 180 msec each. Significant effects of sample colour were discussed earlier; of concern here is the main effect of subinterval, which was significant in 68/174 cells. In 47 of these 68 cases, the significant effect reflected an increase in activity between the first and second subinterval, with the climb sometimes continuing on into the third.

Two examples are shown in Figures 9a and 9b. The usual case is illustrated in Figure 9a; following prolonged inhibition after sample offset, activity climbed through the sample-test interval (200-900 msec post sample onset), but never reached a level greater than pretrial. In Figure 9b, activity at the end of the sample-test interval was substantially above the pretrial level. The result was especially striking for this cell because it was actually unresponsive to the subsequent test display; beyond 900 msec post sample onset the response died away.

Insert Figure 9 about here

Inhibition was even more striking 200-400 msec after offset of the test stimulus. The effect can be seen in the three cells illustrated in Figure 2, and another case is shown in Figure 9c. The effect is certainly not rebound following a positive test response, since it can occur very strongly in cells that

do not respond to the test, and even in cells whose test response is inhibitory. The results suggest a general inhibitory wave entering V4 around 500 msec after a new visual input has been received, perhaps "clearing" the system when processing is complete.

DISCUSSION

Attentional templates

Our findings provide no evidence for the hypothesis of attentional templates in extrastriate cortex. Though targets in our task were defined by colour, and colour specialisation is one of the characteristics of V4, there was no activity either at the time of the sample or in the interval between sample and test that suggested creation or holding of a target description.

A good deal of work would be needed to show why our results differed so strongly from those reported by Haenny et al. (1988) for orientation matching. Their task involved holding a target representation not for one test display as here but through a sequence of displays; their test displays always contained only a single stimulus; the relevant visual dimension was orientation rather than colour, with different numbers of alternatives, discriminabilities etc; and there were doubtless many incidental differences in details of training, level of performance and so on. What our results make clear, however, is that a colour match certainly can be performed with little if any short-term "holding" of the sample in perhaps the most likely candidate area of prestriate cortex.

Earlier we mentioned an alternative to the extrastriate hypothesis: that attentional templates may be carried instead in corresponding areas of the frontal lobe. In light of our findings, this hypothesis seems well worth testing in the current task.

Attentional modulation of the visual response

Before considering their implications, it is worth noting that attentional modulations in this study were really rather weak. Though the bimodality of the distribution in Figure 3 is clear and significant, for very few cells did the direction of attention alter responses by as much as a factor of two, and for most cells the effect was much smaller.

One reason for small effects may simply be that performance was rather poor with a two patch display. It seems likely that, even on trials on which a correct response was made, attention was not always focussed exclusively on the matching patch. At least for the current animal, however, there is no obvious solution since performance is no longer improving. Alternative training methods might include "time-outs" after errors, or presentation of the matching patch slightly in advance of the other patch or patches.

A second important factor could be our use of displays containing at most two patches, since in principle it seems likely that attentional focussing becomes increasingly important as displays contain increasingly large amounts of irrelevant material. To test this idea, we have begun recording from dorsal V4 in a second animal trained with three patch displays. Recordings have been obtained from 13 colour selective cells, of which 9 show a negative attentional modulation and 4 positive. It is too soon to conclude whether effects are stronger than those seen with two patches; the only safe conclusion seems to be that increasing the amount of irrelevant material certainly does not produce a shift towards positive modulation, i.e., towards results resembling those of Moran and Desimone (1985).

Bearing in mind that our effects are rather small, however, there are seven facts that have been established. First, modulation is positive in some

cells but negative in others. Second, the effect whether positive or negative seems extremely local to the cell's preferred region of colour space. Third, modulation is specific to two patch displays; responses to single patches seem largely independent of the preceding sample. Fourth, the effect develops only late in the visual response, at least 100 msec after stimulus onset. Fifth, its size does not depend on the strength of a cell's sensory colour preference. Sixth, response to a two patch display is initially quite close to the response that would be produced by the better of the two single patches it contains; beyond 100 msec, however, a relative inhibition develops for the two patch display, and attentional modulation, whose time course is similar, could perhaps be seen as a relative strengthening or weakening of this general two-patch inhibition. Seventh, attentional modulation is specific for colour, the relevant dimension in this task. There is no evidence for modulation based on location preferences. With these facts in mind we may consider the three alternative hypotheses that were outlined in the introduction.

Input gating

According to the input gating model, attention works by gating out sensory inputs from unattended regions of a cell's receptive field. The findings we have reported are quite inconsistent with this model. First, it predicts that attentional modulation should always be positive: Response should always be stronger when attention is focussed on the cell's more preferred of two stimuli. Second, it predicts a correlation between strength of attentional modulation and strength of sensory preference. Moving attention from one stimulus to another should have an effect that is proportional to their differential effectiveness in stimulating the cell. Third, it predicts that attentional modulations should be the same no matter what dimension is the basis for a cell's preference for one of two stimuli. Though one cell might

give twice as great a response to patch X than to patch Y because of a colour preference, while another gives twice as great a response because of a location preference, in both cases the effect of gating out X vs. Y should be the same. In this study, however, modulations seemed sensitive only to colour, not location preferences.

The input gating hypothesis might well be true under different conditions; for example, when more effective attentional selection is achieved with a spatial cue, as in the study of Moran and Desimone (1985). For the colour matching task, however, this hypothesis is quite inconsistent with the results.

Biassed competition

The biased competition model does no better with these findings. Here the essential idea is that populations of V4 cells with different sensory preferences (e.g., for red vs. green) are mutually inhibitory, and that attention to a given colour is achieved by pre-priming the corresponding population of cells, giving them the edge in subsequent competition. Again there are various features of the results that are inconsistent with such an account. First, in the interval between sample and test we saw no evidence of pre-priming activity. Second, the key prediction of the model is really that cells should always respond more strongly to the test when they have been pre-primed, i.e., when the sample had their preferred colour, and the data in Tables 2 and 3 disconfirm this. For positive cells, for example, response is indeed best given a b sample for a b+w or b+m display, but there is no hint of a similar result with either the m+w or single patch displays, even though the cells are still responding and still presumably competing with other active cell populations that must certainly be firing in all these cases. Unlike the findings of Haenny et al. (1988), our results do not show that cells have a

simple preference for one or another sample, no matter what the test. Instead they suggest something much more elaborate: a specific inhibition that develops either when a b colour patch is ignored (positive cells) or selected (negative cells) in the presence of a second patch or potential attentional competitor.

Feedback modulation

The third hypothesis is that attentional selection takes place outside, perhaps beyond prestriate cortex, with prestriate effects reflecting some sort of feedback input. Though there is little in our findings to support this hypothesis directly, it does seem a promising candidate. It is consistent, for example, with attentional effects developing so late in the response. Perhaps more important, it is logically necessary that selection in our task must be controlled by some sort of advance specification of the target colour; since we found no evidence for such templates in V4, it seems reasonable to conclude that selection must initially be driven by input-template matching at some later stage.

One way to investigate the hypothesis might be to record attentional modulations in V4 while deactivating (e.g., cooling) other structures known to send inputs to it. Perhaps such an experiment could best be saved until direct evidence for attentional templates in one or another such structure had been found.

Positive and negative cells

Perhaps the key issue in formulating an adequate account of our findings is the question of why V4 should have both positive and negative cells. Though we have considered quite a number of hypotheses, we still do not have one that is entirely satisfactory. Broadly speaking, these hypotheses fall into two classes. In the first class are hypotheses proposing that cells are

characteristically positive or negative. For example, they might be anatomically different cells with different functional roles. In the second class are hypotheses proposing that whether a cell appears positive or negative depends on the accident of the task used to assess it.

We have considered at least three hypotheses of the first sort. (a) Perhaps the visual system preserves separate representations of the perceptual "foreground" and "background", used for different purposes. Then a positive cell would be a part of the foreground representation, while a negative cell would be part of the background. (b) Perhaps attention is controlled by some sort of "push-pull" system. Some V4 cells when active tend to draw attention towards the object that is activating them (positive cells), while others (negative cells) tend to push it away. Positive and negative cells responding to the same object (i.e., with the same constellation of sensory preferences) would be mutually inhibitory, and competition within such a network would determine the final focus of attention. (c) Perhaps we are wrong in thinking that a cell always transmits its message more effectively the more strongly it is firing. Instead there is some optimal, intermediate firing rate. Positive cells are those whose baseline response is below the optimal value, while negative cells are those whose baseline response is above it. In either case, attention moves response in the optimal direction.

These hypotheses - and indeed this whole class of hypotheses - share a common difficulty. They are *prima facie* inconsistent with the original results of Moran and Desimone (1985), who using a different selection task found an overwhelming preponderance of positive attentional modulation. If some V4 cells contributed to "background" representations, or had a low optimal firing rate, and so on, then they would show negative attentional modulation no

matter how the focus of attention was controlled. In the last few recording penetrations before animals are sacrificed, we shall place marker lesions near recorded positive and negative cells to test whether they are anatomically distinct, e.g., lying in different cortical layers. On the whole, however, the idea that a given cell is characteristically positive or negative seems unpromising.

It follows that the best way forward may be to find experimental manipulations that change a cell's modulation from positive to negative. One possible lead is suggested by looking again at Figure 5, which shows average responses of all positive and negative cells. Though broadly the two are quite similar, the negative cell response does seem initially to be weaker. Because firing rates are enormously variable across cells, this difference does not approach significance. Nevertheless, might negative cells be those with a weaker preference for our whole set of stimuli, i.e. 1^0 coloured discs? If these cells were tested with other shapes which they preferred, might their attentional modulation turn from negative to positive?

In human vision, the function of attention seems to be to select the visual representation of a whole object; for example, a person asked to read out the single red letter from a display will typically also be able to report the size of this letter, where it occurred and so on (see e.g., Duncan, 1984). Since typically a V4 cell will have multidimensional stimulus preferences - preferences for particular colours, locations, sizes, shapes etc (Desimone & Schein, 1987) - one might suggest that a cell coding some but not all of an attended object's properties receives particularly strong inhibition when the attentional state develops. This would be the multidimensional equivalent of the "Mexican hat" model tested in the six colour study: Attention facilitates a cell whose constellation of preferences is sufficiently close to the properties of

the attended object, but inhibits those whose preferences are just outside this positive zone. Though we can think of difficulties for this hypothesis too, it seems worth testing using stimuli of different shapes.

Another informative manipulation might be to introduce irrelevant stimulus variation along dimensions other than location. We observed attentional modulations based on colour but not location; is this because only colour was relevant for the task, or because coding object locations is not primarily the concern of V4 (Desimone & Ungerleider, 1989)? What would happen, for example, in a colour matching task whose stimuli varied randomly and irrelevantly in shape or orientation?

Whatever the explanation for the existence of both positive and negative cells, our results do make one thing clear. As suggested by the behavioural data, the spatial filtering observed by Moran and Desimone (1985) in V4 is quite different from attentional modulation in a different selection task. We have opened the door to investigation of the general case in which behaviorally relevant visual information is selected on the basis of many different kinds of advance knowledge.

Acknowledgments

Research sponsored by the Air Force Office of Scientific Research, Air Force Systems Command, USAF, under grant number AFOSR-90-0043. Joint funding was provided by the Office of Naval Research. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of the Air Force Office of Scientific Research or the U.S. Government.

Additional financial support was provided by the National Institute of Mental Health, the Medical Research Council, and Award 0085/88 from the NATO Collaborative Research Grants Programme.

References

- Bundesen, C., Shibuya, H. & Larsen, A. (1985) Visual selection from multielement displays: A model for partial report. In: M.I. Posner & O. Marin (Eds.), Attention and performance XI (pp. 631-649). Hillsdale, N.J.: Erlbaum.
- Bushnell, M.C., Goldberg, M.E. & Robinson, D.L. (1981) Behavioural enhancement of visual responses in monkey cerebral cortex. I. Modulation in posterior parietal cortex related to selective visual attention. Journal of Neurophysiology, 46, 755-772.
- Desimone, R. & Schein, S.J. (1987) Visual properties of neurons in area V4 of the macaque: Sensitivity to stimulus form. Journal of Neurophysiology, 57, 835-868.
- Desimone, R. & Ungerleider, L.G. (1989) Neural mechanisms of visual processing in monkeys. In: F. Boller & J. Grafman (Eds.), Handbook of Neuropsychology Vol. 2 (pp. 267-299). Amsterdam: Elsevier.
- Duncan, J. (1983) Perceptual selection based on alphanumeric class: Evidence from partial reports. Perception & Psychophysics, 33, 533-547.
- Duncan, J. (1984) Selective attention and the organization of visual information. Journal of Experimental Psychology: General, 34, 221-232.
- Duncan, J. (1989) Boundary conditions on parallel processing in human vision. Perception, 18, 457-469.
- Duncan, J. & Humphreys, G.W. (1989) Visual search and stimulus similarity. Psychological Review, 96, 433-458.
- Feldman, J.A. (1985) Four frames suffice: A provisional model of vision and space. The Behavioural and Brain Sciences, 8, 265-289.

- Funahashi, S., Bruce, C.J. & Goldman-Rakic, P.S. (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. Journal of Neurophysiology, 61.
- Gnadt, J.W. & Andersen, R.A. (1988) Memory related motor planning activity in posterior parietal cortex of macaque. Experimental Brain Research, 70, 216-220.
- Haenny, P.E., Maunsell, J.H.R. & Schiller, P.H. (1987) State dependent activity in monkey visual cortex: II. Retinal and extraretinal factors in V4. Experimental Brain Research, 69, 225.
- Hillyard, S.A., Munte, T.F. & Neville, H.J. (1985) Visual-spatial attention, orienting, and brain physiology. In: M.I. Posner & O.S.M. Marin (Eds.), Attention and performance XI (pp. 63-84). Hillsdale, N.J.: Erlbaum.
- Merkle, P.M. (1980) Selection from visual persistence by perceptual groups and category membership. Journal of Experimental Psychology: General, 109, 279-295.
- Moran, J. & Desimone, R. (1985) Selective attention gates visual processing in the extrastriate cortex. Science, 229, 782-784.
- Phaf, R.H., van der Heijden, A.H.C. & Hudson, P.T.W. (1990) SLAM: A connectionist model for attention in visual selection tasks. Cognitive Psychology, 22, 273-341.
- von Wright, J.M. (1968) Selection in visual immediate memory. Quarterly Journal of Experimental Psychology, 20, 62-68.
- Wise, S.P. & Desimone, R. (1988) Behavioural neurophysiology: Insights into seeing and grasping. Science, 242, 736-741.
- Zeki, S.M. (1980) The representation of colours in the cerebral cortex. Nature (London), 284, 412-418.

Table 1

Conditions in three colour task

One patch

test display

		red		green		blue	
		x	y	x	y	x	y
sample	r	1	2	3	4	5	6
	g	7	8	9	10	11	12
	b	13	14	15	16	17	18

Two patches

test display

		red + green		red + blue		green + blue	
		r_x	r_y	r_x	r_y	g_x	g_y
		g_y	g_x	b_y	b_x	b_y	b_x
sample	r	19	20	21	22	23	24
	g	25	26	27	28	29	30
	b	31	32	33	34	35	36

Notes: x, y refer to locations

r = red, g = green, b = blue

 r_x = r in location x, etc.

Table 2

Colour selective cells : mean response (impulses/sec) to two patch test displays (130-220 msec post onset)

sample	test display								
	b + w			b + m			m + w		
	b	m	w	b	m	w	b	m	w
positive cells	9.4	8.6	7.3	9.3	7.6	9.1	6.9	7.4	7.1
negative cells	7.1	8.1	9.5	7.0	8.9	8.6	6.5	6.7	6.4

Table 3

Colour selective cells : mean response (impulses/sec) to one patch test displays (130-220 msec post onset)

sample	test display								
	b			m			w		
	b	m	w	b	m	w	b	m	w
positive cells	9.0	8.9	8.7	6.4	6.6	6.5	4.4	4.7	4.7
negative cells	10.2	9.3	9.2	6.2	6.3	6.3	5.0	5.1	4.8

Figure captions

Figure 1. Average response of a V4 neuron following different samples.

Time 0 = sample onset, 200 = offset, 900 = earliest possible onset of test. Solid line - red sample; dotted line - green sample; dashed line - blue sample.

Figure 2. Responses of three neurons to test displays. Histograms show average responses to all one patch displays. Time 0 = test onset, 200 = test offset.

Figure 3. Distribution across cells of attentional modulation index for colour. Only cells meeting the criterion for colour selectivity are included. Data from 130-220 msec post test onset. Extreme bins include all cells outside range ± 1.24 .

Figure 4. Attentional modulation index for colour in the interval 70-130 msec post test onset. Positive and negative cells (as classified by the late response) are separately shown. Extreme bins include all cells outside range ± 1.24 .

Figure 5. Average response of positive and negative cells to displays containing the b colour. Dashed line - single b patch, average for all samples. Solid line - average for b+w and b+m displays, sample = b. Dotted line - average for b+w (sample = w) and b+m (sample = m) displays. Only cells recorded with patch separation = 45° are included ($N = 32$ positive, 34 negative cells). Time 0 = test onset, 200 = test offset.

Figure 6. Distribution across cells of attentional modulation index for location. Only cells meeting the criterion for location preference are included. Data from 130-220 msec post test onset. Extreme bins include all cells outside range ± 1.24 .

Figure 7. Alternative models of attentional modulation. (a) Hypothetical sensory tuning curve. (b) Model 1: Separate attentional inputs for positive

and negative cells. Model 2: All cells have both positive and negative attentional inputs from different regions of colour space.

Figure 8. Mean response (40-220 msec post test onset) to single patches of different colours. (a) Cell with typical "Mexican hat" tuning curve. (b) Cell with stronger upturn for colour opposite to the best. (c) Average tuning curve for all colour selective cells. Best = adjacent pair of colours giving strongest average response; far = opposite pair to best on colour circle; middle = intermediate colours.

Figure 9. (a) Average response of a V4 neuron following different samples. Time 0 = sample onset, 200 = offset, 900 = earliest possible onset of test. Solid line - red sample; dotted line - green sample; dashed line - blue sample. (b) As for (a); a different neuron. (c) Average response of a V4 neuron following different tests. Time 0 = test onset, 200 = offset, 900 = earliest possible onset of next stimulus (on mismatch trials). Solid line - red test; dotted line - green test; dashed line - blue test.

Figure 1






Figure 2


(a) ch-121-2.0(1)



(b) ch-123-2.0(1)



(c) ch-124-1.0(2)



Colour data
130-220 msec post onset

Figure 3





Figure 4

Colour data
70-130 msec post onset

(a) Positive cells



(b) Negative cells




Figure 5




Figure 6






Figure 7

(a) Sensory input



(b) Attentional input

Model 1



Model 2





Figure 8

(a) ch-124-3.0 (1)



(b) ch-162-1.0(1)



(c) Mean






Figure 9

(a) che-93-1.0(1)



(b) che-20-1.0(1)



(c) ch-132-1.0(1)

