

Attention Modulates the Responses of Simple Cells in Monkey Primary Visual Cortex

Carrie J. McAdams and R. Clay Reid

Department of Neurobiology, Harvard Medical School, Boston, Massachusetts 02115

Spatial attention has long been postulated to act as a spotlight that increases the salience of visual stimuli at the attended location. We examined the effects of attention on the receptive fields of simple cells in primary visual cortex (V1) by training macaque monkeys to perform a task with two modes. In the attended mode, the stimuli relevant to the animal's task overlay the receptive field of the neuron being recorded. In the unattended mode, the animal was cued to attend to stimuli outside the receptive field of that neuron. The relevant stimulus, a colored pixel, was briefly presented within a white-noise stimulus, a flickering grid of black and white pixels. The receptive fields of the neurons were mapped by correlating spikes with the white-noise stimulus in both attended and unattended modes. We found that attention could cause significant modulation of the visually evoked response despite an absence of significant effects on the overall firing rates. On further examination of the relationship between the strength of the visual stimulation and the firing rate, we found that attention appears to cause multiplicative scaling of the visually evoked responses of simple cells, demonstrating that attention reaches back to the initial stages of visual cortical processing.

Key words: attention; striate; vision; cortex; macaque; mapping; receptive field

Introduction

Attention increases the responses of neurons in many areas of visual cortex (for review, see Treue, 2001). However, attentional modulation of the responses of neurons in primary visual cortex (area V1) in monkeys has been inconsistently observed, despite robust attentional effects using the same tasks in extrastriate areas (Motter, 1993; Luck et al., 1997; Roelfsema et al., 1998; Vidyasagar, 1998; Ito and Gilbert, 1999; McAdams and Maunsell, 1999; Grunewald et al., 2002; Marcus and Van Essen, 2002). The strongest reported effects of attention in V1 have been found in the responses to thin lines placed within V1 receptive fields (Motter, 1993; Roelfsema et al., 1998; Ito and Gilbert, 1999). Furthermore, the largest enhancements occurred either late in a long stimulus presentation period (Roelfsema et al., 1998) or required the presence of additional stimuli near the receptive field of the neuron (Ito and Gilbert, 1999). Other studies have found that attention affects V1 little (McAdams and Maunsell, 1999) or not at all (Luck et al., 1997; Grunewald et al., 2002; Marcus and Van Essen, 2002), but most of these studies used stimuli that were much larger than V1 receptive fields.

Previous examples of attentional effects in V1 were found in either superficial layer complex cells (Ito and Gilbert, 1999) or cells of unspecified type (Motter, 1993; Roelfsema et al., 1998; McAdams and Maunsell, 1999). We wanted to examine whether

attention modulates the responses of V1 simple cells, whose receptive fields are generally smaller than V1 complex cells and which are arguably earlier in the cortical processing stream. We used a task that required the discrimination of small, briefly presented stimuli on a noisy background in the hope that it would require the information carried by simple cells and might therefore offer the best situation for observing attentional modulation of their responses.

Spatial attention is often likened to a spotlight that can improve the detection and discrimination of visual stimuli (Posner et al., 1980; Hurlbert and Poggio, 1985; Humphreys and Bruce, 1989; Cave and Bichot, 1999). Neurophysiologically, attention has also been shown to produce small improvements in stimulus discriminability without changing the underlying stimulus selectivity of the neurons for orientation (McAdams and Maunsell, 1999) and direction (Treue and Martinez Trujillo, 1999; Recanzone and Wurtz, 2000), consistent with the spotlight metaphor. Other studies have shown that the effective weighting of neuronal receptive fields can shift depending on the locus of attention (Moran and Desimone, 1985; Connor et al., 1996, 1997; Treue and Maunsell, 1999; Cook and Maunsell, 2004), yielding enhanced representation of attended objects. Here, we mapped the receptive fields of simple cells, with and without spatial attention directed to the receptive field. We found that attention enhanced the visual responses of simple cells in primary visual cortex, without changing the underlying spatial or temporal structure of the receptive fields.

Materials and Methods

Animal care and the experimental protocol were in accordance with the National Institutes of Health and United States Department of Agricul-

Received March 22, 2005; revised Oct. 5, 2005; accepted Oct. 19, 2005.

This work was supported by National Institutes of Health Grants R01 EY12815, P30 12196, and T32 NS07484. We thank S. Yurgenson for computer support, J. Reppas and J. Maunsell for comments on this manuscript, and J. Assad for helpful discussions and assistance.

Correspondence should be addressed to Carrie J. McAdams, Department of Neuroscience, Harvard Medical School, 200 Longwood Avenue, Boston, MA 02115. E-mail: cmcadams@hms.harvard.edu.

DOI:10.1523/JNEUROSCI.2904-05.2005

Copyright © 2005 Society for Neuroscience 0270-6474/05/2511023-11\$15.00/0

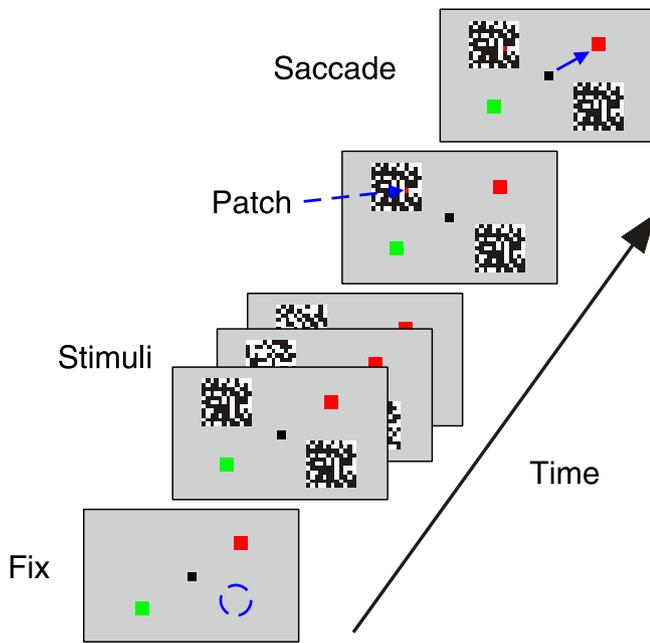


Figure 1. Stimulus configuration and behavioral protocol. The animal fixated on the central spot and had previously been cued to direct his attention to one of the two stimulus grids. The blue dashed circle represents the receptive field of the recorded neuron and was not present on the screen. Trials in which the animal's attention was directed to the grid overlying the receptive field were defined as attended, whereas those trials in which his attention was directed to the other grid were defined as unattended. A small colored patch (in this example, follow the dashed blue arrow to a red square in the top left stimulus grid) was visible for 250–350 ms, during which time all of the other pixels in the stimulus continued to change luminance randomly. If the animal's attention had been directed to the grid on the top left, the correct response would be a saccade to the red target, illustrated by the solid blue arrow. If the animal's attention had been directed to the stimulus grid on the bottom right (no colored patch shown in this example), the correct response would have been to maintain fixation. Such catch trials ended after a random interval ranging from 260 to 3300 ms.

ture guidelines and were approved by the Harvard Medical Area Standing Committee on Animals.

Behavioral task and training. Two macaques (one *Macaca mulatta* and one *Macaca fascicularis*) were used. Each animal had a scleral search coil inserted in one eye and a head post affixed to the skull during an aseptic surgery before training started. The animals were rewarded with fruit juice or water for performing the desired behaviors. One animal was trained for 11 months and the other for 7 months before data collection began.

The animals did a task in which two regions of white-noise stimuli appeared. One region was cued as the relevant stimulus at the start of each trial. Two small colored spots, the “targets,” were also present on the screen throughout each trial. The animals were trained to detect the occurrence of a colored pixel, the “patch,” within the cued stimulus area and immediately saccade to the target that had the same color as the patch and to continue fixating if no color patch appeared in that stimulus (Fig. 1). They had one to three instruction trials in which only one white-noise stimulus was present before the two stimuli were shown. Additional cueing trials were provided if more than three trials in a row were incorrectly completed or ignored. Colored pixels also appeared in the uncued stimulus, but animals were not rewarded for responses to these. In each block and for each of the two stimulus locations, one-third of trials had a green patch, one-third of trials had a red patch, and one-third of trials had no colored patch. Two trials of each possible combination (red patch on stimulus 1 and 2, red patch on stimulus 1 and green on 2, red patch on 1 and no patch on 2, etc.) were collected before switching the stimulus to which the animal was directing his attention. Each block of data consisted of 18 trials during which the animal's attention was directed to one stimulus location and 18 trials during which the animal's attention was directed to the other stimulus location. Because only one stimulus was

over the receptive field of the neuron being recorded, the trials in which the animal attended to that stimulus are referred to as the “attended mode,” and the trials in which the animal attended to the other stimulus are referred to as the “unattended mode.” Thus, throughout the paper, attended and unattended refer to the location of the animal's attention relative to the receptive field of neuron that we are studying. The time at which each patch might appear in each stimulus was randomly and independently selected from a minimum (60–100 ms) to maximum (2400–2900 ms) after the start of the stimulus sequence. The minimum and maximum values were fixed for each cell. Only correctly completed trials were included in the neuronal data analysis.

Stimulus. The stimulus grids, arrays of 8×8 to 12×12 black and white pixels, were presented on a video monitor with a video refresh rate that was fixed for each recording day and set to either 65 or 85 Hz. Each pixel was a 0.2 – 0.3° square whose luminance was determined by a pseudorandom binary temporal signal (Reid et al., 1997). This white-noise stimulus was updated every two or three video frames, also fixed for each recording day. Thus, a single stimulus frame of white noise was shown for a duration of 23.5, 30.5, or 34.5 ms, depending on both the video refresh rate and stimulus frame refresh rate. The color patches were randomly assigned to any one of the pixels within the stimulus on each trial and were shown for 250–350 ms. We chose the red and green color values such that the animals performed at similar levels on the task for both patch colors. The resulting red and green patch colors were not photometrically equiluminant.

We attempted to optimally size and position the stimulus for each neuron. We excluded from additional analysis neurons that were not driven by the stimulus, those that did not show any structure in their receptive field based on the white-noise mapping procedure, and those whose receptive fields did not appear to be simple. Recordings made from simple cells whose receptive fields were much smaller or much larger than the pixels of the stimulus would not show the structure of a simple cell receptive field and therefore have also been excluded. These results are thus biased to the population of simple cells whose receptive fields were at a minimum size of two of the pixels we used (0.2 – 0.3° , one overlying each subfield). We used the smallest pixels at which the animal was able to perform the task accurately at a given eccentricity, unless the receptive field exceeded the overall stimulus size (2.4 – 3.6° square) or the neuron only responded to larger pixels.

We recorded complete data sets from 138 neurons, 46 of which ultimately had receptive fields meeting our criteria for simple cells, 10 of 26 from animal 1 (*M. mulatta*) and 36 of 112 from animal 2 (*M. fascicularis*). The average receptive-field size of these simple cells ($\sim 1.0^\circ$ per subfield at 4° eccentricity) is larger than that reported previously for V1 neurons in anesthetized macaques (Hubel and Wiesel, 1974; Dow et al., 1981; Van Essen et al., 1984). This is most likely attributable to the large pixels in the visual stimuli, which might have resulted in both an overestimate of receptive-field size and a selection bias for neurons with larger receptive fields.

Recording. After the animal learned the behavioral task, a recording chamber was positioned over primary visual cortex in an aseptic surgery. The bone was left intact, and small craniotomies were drilled inside the chamber as needed to provide access to the cortex. The dura was left intact and penetrated daily using homemade glass-coated platinum/iridium electrodes of impedance ranging from 1 to 4 M Ω at 1 kHz. Each craniotomy provided access for between 2 and 8 weeks of data collection.

While the animal performed the behavioral task, single neurons were isolated as the electrode advanced. When a unit was isolated, data collection was initiated. All units were isolated within 1800 μm of the cortical surface and before the electrode encountered any stretches of white matter. Between 5 and 14 blocks of the behavioral task were used in the analysis, providing an average of ~ 5000 frames of the white-noise stimuli and ~ 3200 spikes.

The spike waveforms were sampled at 11 kHz and the animal's horizontal and vertical eye positions at 900 Hz. The mean spike rate in response to the white-noise stimulus was highly variable across cells (range of 0.5–117 spikes/s). The spike waveforms were analyzed off-line using Spike 2 (Cambridge Electronic Design, Cambridge, UK) to ensure isolation of single units. In two cases, distinct waveforms representing two

simple cells, of different orientation or receptive-field size, were recorded simultaneously on one electrode, and later both were isolated and included in this analysis.

Analysis. Spatiotemporal receptive fields (or linear kernels) were mapped by correlating the white-noise stimulus with spikes that followed it at multiple time delays, using 10 ms bins (Reid et al., 1997). This calculation yields the average firing rate of the neuron, above or below the mean, after the bright phase of the stimulus at each pixel, calculated for each delay, or time bin, between stimulus and response. Only those spikes that occurred at least 100 ms after the initial stimulus onset and before the appearance of a color patch at either stimulus position were used in creating the spatiotemporal receptive-field maps. Simple cells have a characteristic receptive-field structure consisting of at least two parallel, antagonistic subfields (Hubel and Wiesel, 1962, 1968; Schiller et al., 1976). We defined neurons as simple cells if their receptive fields had two or more spatially distinct regions of opposite polarity that were not center-surround.

We extracted the spatial receptive field and the time course of the response from the spatiotemporal receptive field in a multistep procedure. First, the time bins with the strongest responses were determined by summing the absolute values of all pixels in each time bin. Next, we took the one to three time bins with the strongest responses (greater than half the maximum) and added them together for each pixel to yield what we will term the “spatial receptive field.” The “primary subfield” was defined as the pixels in the spatial receptive field that had responses >1 SD of the noise (defined below), were of the same polarity as the strongest pixel, and were contiguous with that pixel and others meeting those criteria. The “secondary subfield” was defined similarly but for the pixels of the opposite polarity. The primary and secondary subfields were defined independently for each behavioral mode, and only those pixels that overlapped across the two modes were used in additional comparisons. Finally, the “time courses” of the primary and secondary subfields were calculated by summing over all pixels in a subfield for each time bin.

The noise in our measurements of the receptive fields depended on the number of stimulus frames shown, the number of spikes, and the number of pixels composing the receptive field. We quantified the noise with two methods: the “baseline noise” and the “visual response noise.” The baseline noise was the SD of the response values for pixels at time bins outside the stimulus-elicited response (110–160 ms). For a neuron whose response can be approximated by a Poisson process, the visual response noise should be well estimated with this baseline noise. We confirmed this by computing the visual response noise using a Monte-Carlo simulation. For each of 128 iterations, we randomly sampled half of the spikes to calculate 128 different spatiotemporal receptive fields. We then determined the SDs of bins with the largest responses (those used to calculate the total visual response; see below). The average of these SDs were defined as the visual response noise. We found that the baseline noise and the visual response noise were nearly equivalent for most neurons (paired *t* test, $p = 0.08$; mean baseline noise, 0.66 spikes/s; mean visual response noise, 0.64 spikes/s). Half of the neurons showed a $<5\%$ difference in the two measures of noise and only five had a $>20\%$ difference. Finally, there were no neurons for which using the baseline noise as opposed to the visual response noise changed the significance of the responses. We therefore used the baseline noise in all calculations of significance.

In our analysis of the statistical significance of attentional modulation, we considered first the time course for each subfield. If the summed responses from pixels in a 10 ms bin of the time course were more than three times the estimated noise [the single-pixel SD (baseline noise) times the square root of the number of pixels contained in the subfield], we included that time bin in the calculation of significance. Typically, neurons had visual responses that met these criteria for 20–80 ms, so the temporal summation included two to eight time bins, with a median of five. We defined the “total visual response” of a subfield as the sum of the absolute values of each time bin that met this criterion. These time bins are indicated with asterisks along the *x*-axis in the figures. The difference in the total visual response for the two behavioral modes had to exceed three times the noise of our measurement of the total visual response (baseline noise multiplied by the square root of the total number of pixels

summed over both space and time) for a neuron to be classified as having an individually statistically significant effect of attention. This threshold for single time bins (baseline noise times the square root of the total number of pixels in the receptive field summed over space) is illustrated with light gray lines in each figure showing the time course of a visual response. A difference at a single time bin would be considered statistically significant if the difference between the attended and unattended curves exceeded this threshold. For the total visual response of a cell to be significant, the sum of the differences across the selected time bins had to exceed this value times the square root of the number of time bins included in the total visual response.

All statistical tests on the visual response were performed on the raw data, the actual values of the summed pixels, although the impulse-response functions and some of the spatiotemporal receptive-field maps may be shown as smoothed curves in some of the figures. For our eye position control, the actual measured eye position was used to offset the white-noise stimuli on a frame-by-frame basis, to a quarter pixel resolution. The receptive-field subregions were then defined in the same way as previously, except at quarter-pixel resolution. The reported values and statistical tests on the spike rates were performed on the actual driven rates during visual stimulation, without any correction for the undriven activity of the neuron.

Analysis of the responses in terms of a two-stage model, a linear pre-filter followed by an output nonlinearity (see Figs. 9, 10), was based on methods described in detail by Chichilnisky (2001). Such a model allows for a simple analysis of the output nonlinearity (Hunter and Korenberg, 1986), although it is only an approximation (Victor, 1992). We first created a linear kernel (spatiotemporal receptive field) for each neuron using the data from all trials. Based on their similarity to the linear kernel, each stimulus configuration was entered into a particular linear input bin, along with the corresponding number of spikes it evoked (actual firing). Because attention did not systematically alter the spatial or temporal structure of receptive fields (see Fig. 7), the same linear first stage was used to analyze both attended and unattended trials. The output of the linear stage was then estimated on a frame-by-frame basis by convolving the kernel with the visual stimulus (see Fig. 9). The output nonlinearity was estimated by comparing the result of the linear stage with the actual firing of the neuron (see Fig. 10).

Results

Behavior

Two macaque monkeys found a single colored pixel (the patch) that appeared among noise: a grid of small ($0.2\text{--}0.3^\circ$ at $3\text{--}5^\circ$ eccentricity), randomly flickering black and white pixels (Fig. 1) (see Materials and Methods). The task was difficult for the animals, whose average performance was 76%. Across all cells and excluding trials with fixation breaks before any patch appeared (28%), the animals' overall performance on the task was 75% correct when their attention was directed to the stimulus grid over receptive field and 77% when their attention was directed to the other grid. Most errors were failures to respond (miss, 19% of responses) rather than responding erroneously (false alarm, 3.5% of responses). The two animals had different error patterns. Animal 1 missed the target on 38% of the saccade trials and gave false alarms on 18% of the fixation trials. Animal 2 missed the target on 22% of saccade trials and gave false alarms on 28% of the fixation trials. Both animals were slightly faster in responding to the red patch than to the green patch (mean response time, 410 vs 447 ms).

Because the animal was always asked to direct all of its attention to one of the two stimuli, we cannot use the animal's performance to estimate changes in the animal's ability to detect or discriminate the colored patch based on the locus of the animal's attention. However, there were some trials in which the uncued patch preceded the cued patch and was also the opposite color (“uncued-patch first trials”). If the animal then saccaded to the color target that matched the color of the uncued patch, this

indicated that the animal's attention was not effectively focused on the cued stimuli. This type of error provides an estimate of how well the animal's attention was focused on the correct stimulus. We can compare this rate of response with the overall patch response rate from the saccade trials. Animal 1 rarely made this type of error, responding to patches in the uncued stimulus an average of only 2% of the uncued-patch first trials and making correct saccades in 56% of the saccade trials. Animal 2 more frequently made this type of error, responding to the uncued stimulus on an average of 17% of the uncued-patch first trials and making correct saccades on 71% of the saccade trials. The ratio of the correct saccade performance on saccade trials to the wrong saccades on uncued-patch first trials is a measure of how well the animal is attending to the cue and will approach 1 as the animal ignores or forget the cue. Animal 1 had an average behavioral ratio of 39, and animal 2 had an average behavioral ratio of 5. These values suggest that both animals were aware of the cue and tried to selectively focus their spatial attention on only one of the two stimuli on any given trial.

Receptive-field structure

Receptive fields were mapped by correlating pixel values of the white-noise stimulus (white pixels, 1; black pixels, -1) with spikes that followed them at multiple time delays (Reid et al., 1997). These spatial and temporal maps are in units of spikes per second, and we will refer to them as the visual response. The visual response is a measure of the correlation of the firing of the neuron with the pixels of the stimulus and is not the actual firing rate. Rather, it is the modulation in the firing rate that would be predicted if a pixel displayed a particular luminance at a specific time. One set of maps was obtained while the animal attended to the stimulus grid overlying the receptive field (the attended mode); the other was obtained while the animal attended to the grid outside the receptive field (the unattended mode). For the simple cell illustrated in Figure 2, the spikes were most strongly correlated with stimuli that had been presented 23.5–47 ms earlier (two frames). This cell has biphasic responses in both subfields. The primary (stronger) subfield was initially off (spikes in response to black pixels), indicated by the blue regions of the receptive-field map; the secondary (weaker) subfield was initially on (spikes in response to white pixels), shown in red. The responses changed sign, or rebounded between 47 and 70.5 ms (the third frame). At all latencies, this neuron had stronger responses in the attended mode (Fig. 2A, bottom) than in the unattended mode (Fig. 2A, top), although the overall firing rate was barely changed (mean response rate over all correct trials, unattended mode, 10.4 spikes/s; attended mode, 10.6 spikes/s).

To examine the amplitude and time course of the receptive-field maps more quantitatively, we summed over all pixels in the

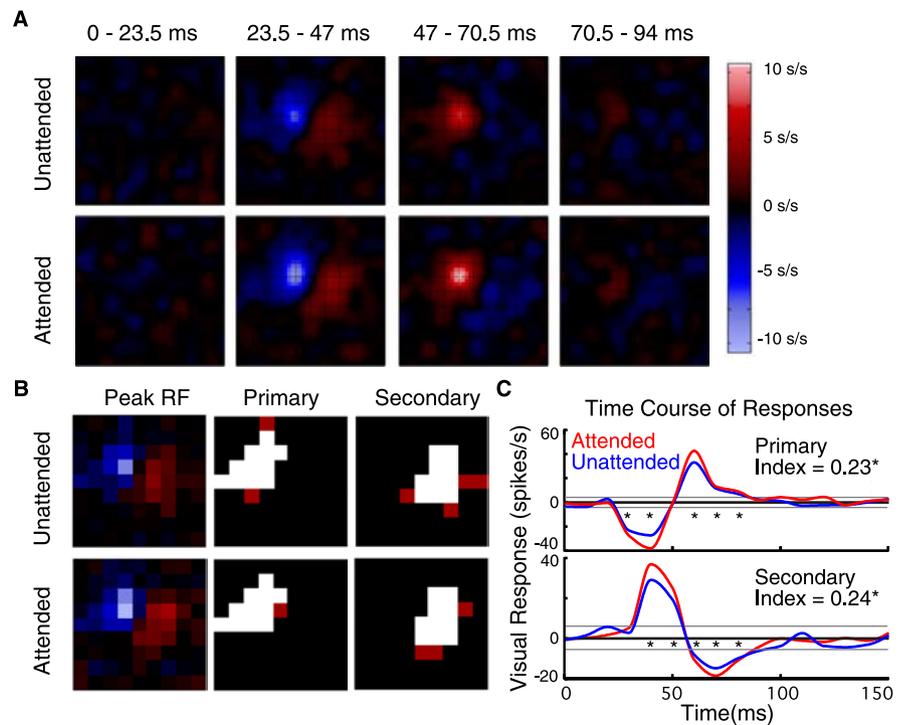


Figure 2. *A*, The spatiotemporal receptive-field map of a single neuron in the unattended mode (top) and attended mode (bottom). Blue indicates regions in which black pixels tended to precede spikes (off responses, negative values), and red indicates regions in which white pixels preceded spikes (on responses, positive values). Brighter colors indicate more responsive areas of the receptive field. Intersections in the grid lines correspond to centers of the stimulus pixels. Response values were interpolated between pixels. The unattended maps were generated from 2473 spikes and 5260 frames over 108 correct trials, and the attended maps were generated from 2438 spikes and 5160 frames over 108 correct trials. *B*, The spatial receptive-field (RF) maps were extracted from the spatiotemporal maps by averaging the time frames near the peak response and selecting pixels meeting threshold levels and contiguous with other strong pixels (see Materials and Methods). This was done independently for the primary (column 2) and secondary (column 3) subfields and for both behavioral modes (red and white pixels). However, only those pixels that overlapped in the two behavioral modes were included in the comparison (white pixels). *C*, The time courses of the responses for the receptive-field subfields, summed over all white pixels defined in *B*. Attended modes are shown in red, and unattended modes are in blue. The gray lines show ± 3 SDs of the noise (pixel variability times the square root of the number of pixels composing the spatial receptive-field subregion) in the response. The asterisks along the x-axis indicate those time bins in which the responses were > 3 SDs of the noise for both the attended and unattended modes and were summed to determine the total visual responses for each mode and each subregion.

primary and the secondary subfields (see Materials and Methods) (Fig. 2B) and plotted the summed response as a function of time, using 10 ms bins (Fig. 2C). For this neuron, attention resulted in a significant 25% increase in the total visual response (see Materials and Methods) of the primary subfield and a significant 28% increase in the total visual response of the secondary subfield. The “attentional index” was defined as the total visual response in a subfield (see Materials and Methods) in the attended mode minus the total visual response in the unattended mode divided by their mean. The attentional index value for this neuron for the primary subfield was 0.23, and the attentional index value for the secondary subfield was 0.24. The behavioral ratio was 5.1 while this cell was recorded (see Results, Behavior), which indicates that this animal (animal 2) was attending to the cued patch in performing the task.

Most of the neurons had smaller effects of attention. The receptive fields and time courses of the subfields for two additional neurons are shown in Figure 3. The neuron in Figure 3A showed no significant effects of attention but had an 11% increase in the total visual response of its primary subfield and a 19% increase in the total visual response of its secondary subfield. This cell had only a minimal change in firing rate: unattended, 7.0 spikes/s; and

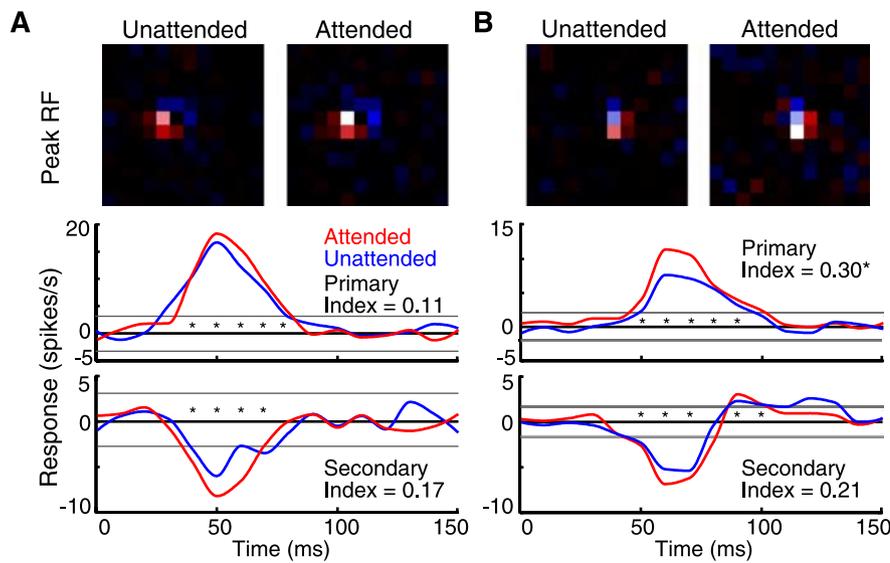


Figure 3. Spatiotemporal receptive-field maps (unattended, left; attended, right) and time courses of the responses of the receptive-field (RF) subfields (primary, top; secondary, bottom) for two single neurons that showed modest effects of attention. For the neuron in **A**, the unattended maps were generated from 1370 spikes and 4573 frames over 108 correct trials, and the attended maps were generated from 1559 spikes and 4991 frames over 108 correct trials. For the neuron in **B**, the unattended maps were generated from 2170 spikes and 5924 frames over 144 correct trials, and the attended maps were generated from 2539 spikes and 5526 frames over 144 correct trials.

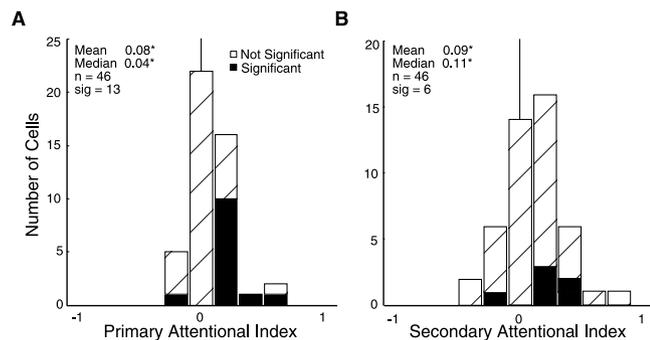


Figure 4. Histograms of the attentional modulation index for the primary subfield (**A**) and the secondary subfield (**B**). The attentional modulation index was defined as the attended response minus the unattended response divided by their mean. Cells shown in black had individually statistically significant effects of attention.

attended, 7.3 spikes/s. The behavioral ratio for animal 2 while this cell was recorded was 3.2. The neuron in Figure 3B had a significant 35% increase in the total visual response of its primary subfield and an insignificant 23% increase in the total visual response of its secondary subfield. Attention also increased the firing rate of this neuron by 27% (attended, 12.4 spikes/s; unattended, 9.8 spikes/s). The behavioral ratio for animal 1 while these data were collected was 49.4.

We recorded from 46 simple cells in the primary visual cortex of two macaques. Within this sample, there was a significant increase in the responses associated with the receptive field in the attended mode for both the primary subfields (Fig. 4A) (mean attentional index, 0.08; *t* test, $p < 0.02$) and the secondary subfields (Fig. 4B) (mean attentional index, 0.09; *t* test, $p < 0.03$). Another interesting question is whether the behavioral ratio, our estimate of how well the animal was attending to the correct stimulus, was correlated with the attentional index, our measure of the change in the receptive-field maps of the neuron. We found that there was no significant correlation between the neuronal

attentional index and the animal's behavioral ratio for either animal (animal 1, $r^2 = 0.08$, $p > 0.5$; animal 2, $r^2 = 0.02$, $p > 0.4$).

Because receptive fields in V1 are small, great care must be taken that changes in eye position do not affect the measurement of the visual responses in the alert macaque. We assessed this concern in several ways. First, we observed no systematic effects on the distribution of eye positions and their variance across the two behavioral modes for both animals (Table 1). The mean differences in eye position were $< 1/20$ of the smallest pixels used, and there was no difference in the variance in the eye positions. The mean pixel size was 0.22° , and the average size of the primary subfield (white pixels only as in Fig. 2B) was 4.7 pixels. Second, because our task actually compared the receptive-field maps under the two conditions, a consistent difference in eye position would be expected to shift the receptive-field maps, not change the measured response amplitudes. A pronounced receptive-field shift was never observed, nor is it likely that

such a shift would cause a systematic offset that consistently favored the map in the attended mode.

As an additional control, we subdivided the actual measured eye position during each stimulus frame into quarters of the stimulus pixels and then created receptive-field maps based on the measured eye position within the fixation window. We reconstructed the receptive-field maps based on the measured eye position for 45 cells. Two examples of eye position reconstructed receptive fields are shown in Figure 5. We found that the corrected receptive fields were slightly smaller (mean primary subfield size, before correction, 4.7 pixels; after correction, 4.2 pixels; paired *t* test, $p < 0.02$). This suggests that small eye movements blur the white-noise stimulus resulting in larger receptive-field maps and that our measurement of eye position within the window increases the accuracy of our receptive-field measurements. More importantly, there were no substantial changes in the attentional modulation of the isolated primary subfield after correction for eye position (Fig. 6A) ($n = 45$; mean index value, 0.08; *t* test, $p < 0.02$). In Figure 6B, the attentional modulation index of the primary subfield without compensation is plotted against the attentional modulation index for the same subfield with eye position compensation. The two are highly correlated ($r^2 = 0.69$; $p < 0.0001$), suggesting that changes in eye position did not systematically affect our results.

An important question is whether the observed attentional modulation changed spatial or temporal aspects of the visual responses or simply modulated the response strength. Attention did not alter the time of the peak response (median time-to-peak, unattended, 60 ms; attended, 60 ms; Wilcoxon's signed rank, $p = 0.96$) or the size of the receptive fields (red and white pixels of primary subfield as in Fig. 2B) (median, 5.8 pixels unattended; 5.8 pixels attended; Wilcoxon's signed rank test, $p = 0.31$). We further considered the temporal aspects of the visual responses by normalizing and averaging the time courses of the primary subfields (as in Fig. 2C) over the 12 neurons with significant positive effects of attention. The curves for the attended responses, unat-

Table 1. Eye position changes across behavioral mode for each animal

	Horizontal eye position					Vertical eye position				
	Difference		SD			Difference		SD		
	Mean	<i>p</i>	A	U	<i>p</i>	Mean	<i>p</i>	A	U	<i>p</i>
Animal 1	0.005°	0.74	0.21°	0.21°	0.79	0.01°	0.40	0.12°	0.12°	0.52
Animal 2	0.006°	0.09	0.11°	0.11°	0.91	0.01°	0.15	0.08°	0.08°	0.29

The mean difference in eye positions (attended – unattended) and the mean SDs for the attended (A) and unattended (U) modes are shown in degrees. The *p* values were determined from a paired *t* test on the mean eye positions or the mean SDs for the attended and unattended modes for all neurons from each animal.

tended responses, and their differences all had essentially the same shape (Fig. 7A). Furthermore, the averaged differences were significantly different from zero starting before the peak of the temporal response curves, showing that attentional modulation can be present even in the early phase of a visually evoked response. The most salient spatial components of the receptive field of a simple cell are the antagonistic spatial subfields. One might expect them to be differentially modulated by attention to the extent that the subfields may derive from different mechanisms (excitatory vs inhibitory, or feedforward vs recurrent). Individual neurons tended to show similar amounts of attentional modulation on both their primary and secondary subfields (Fig. 7B) ($r^2 = 0.28$; $p < 0.001$). Therefore, both the temporal and spatial components of the receptive field show increased response strength without changes in receptive-field structure.

Firing rate

In the analyses above, we examined the effect of attention on the visually driven response, which depends on the degree of the correlation between the spikes and the visual stimulus. A different question is whether the mean spike rate of the neurons was altered by attention. The white-noise stimulus excited most neurons, with only 3 of 46 neurons showing suppressed firing rates in response to stimulation. The range of mean firing rates in response to the white noise varied from 0.5 to 115 spikes/s, with a median of 10.5 spikes/s and a mean of 16.3 spikes/s. On average, we found that attention caused minimal changes in spike rate, during both the presentation of the white-noise stimulus [mean spike rate during stimulation, unattended, 16.3 spikes/s; attended, 16.3 spikes/s; paired *t* test, $p = 0.85$ (Fig. 8A); mean index value, 0.04; *t* test, $p = 0.05$] and the period preceding stimulus onset when the receptive field contained only a gray screen (mean undriven spike rate, unattended, 5.8 spikes/s; attended, 5.6 spikes/s; paired *t* test, $p = 0.09$). We also found no correlation between the magnitude of the overall firing rate in response to the white-noise stimulus and the attentional modulation of the receptive-field map (attentional index vs mean spike rate, $r^2 = 0.01$; $p = 0.50$). However, the population showed a modest correlation between the attentional modulation of the mean spike rate (attended spike rate minus unattended spike rate divided by the mean spike rate) and the attentional modulation of the visually driven response (Fig. 8B) ($r^2 = 0.25$; $p < 0.001$).

Although a linear analysis captures the main features of the

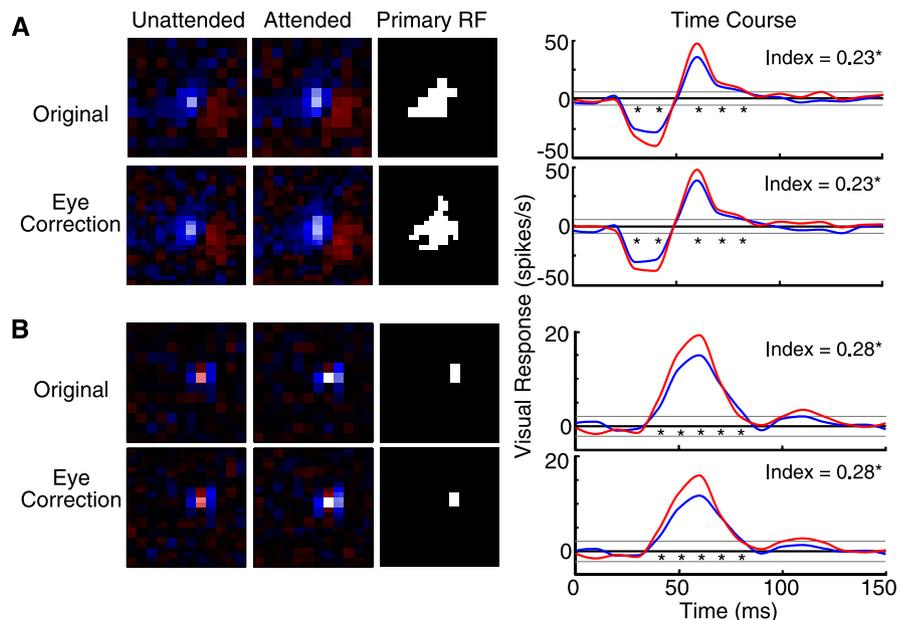


Figure 5. Spatial and temporal receptive fields calculated with and without correction for eye position. Two different neurons are shown in **A** and **B**. For each, the spatial receptive fields (RF) in the unattended, attended, and the pixels chosen for analysis of the primary subfield are shown on the left. The time course of the responses in those pixels is shown on the right. The top row for each neuron are the original (uncorrected) maps, and the bottom row are the maps after adjusting for the actual eye position relative to the white-noise stimulus, at a resolution of a quarter of each pixel, at the time of the spikes. Both neurons had individually significant effects of attention both with and without eye correction. Neuron A is the same neuron shown in Figure 2. The attentional index value for each condition is displayed on the top left of the time response function.

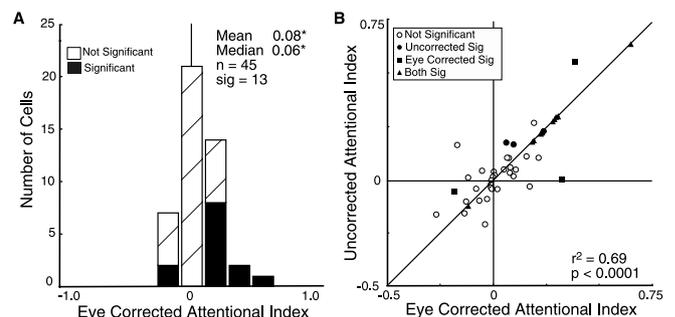


Figure 6. **A**, Histogram of the attentional modulation index of the primary subfield (as in Fig. 4), calculated with the correction for eye position. The 13 cells shown in black had individually statistically significant effects of attention. **B**, The attentional index for the primary subfield without the correction for eye position is plotted against the attentional index with the correction for eye position. The diagonal $x = y$ is plotted as well. Filled symbols indicate cells with significant effects: filled circles indicate effects only without eye correction (3 cells); filled squares indicate effects only after eye correction (3 cells); and filled triangles indicate effects in both cases (10 cells).

receptive field of a simple cell, the responses of simple cells are not simply linear transformations of their visual inputs. Most notably, their responses are rectified: increasing levels of depolarization lead to increasing probabilities of firing, but hyperpolariza-

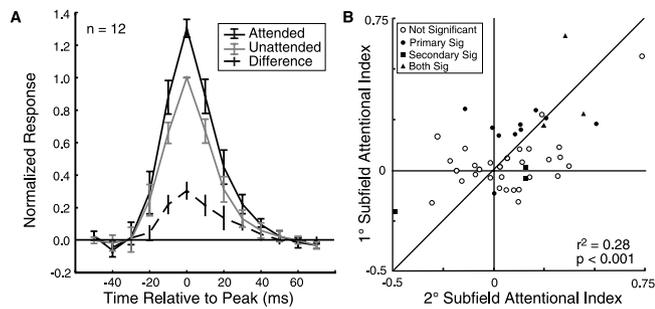


Figure 7. *A*, The time course of the responses for the 12 neurons with individually significant positive effects of attention were constructed by normalizing each cell to its peak response in the unattended mode and then shifting each function so that the peak responses of each neuron were aligned in time. The attended mode is in black, the unattended mode in dark gray, and their difference is shown with a dashed black line. *B*, The attentional index for the peak of the primary subfield is plotted against the attentional index for the peak of the antagonistic subfield. The diagonal $x = y$ is plotted as well. Filled symbols indicate cells with significant differences between the attentional states: filled circles indicate an effect on the primary subfield only (10 cells); filled squares indicate an effect on the secondary subfield only (3 cells); and filled triangles indicate the cells with effects on both subfields (3 cells).

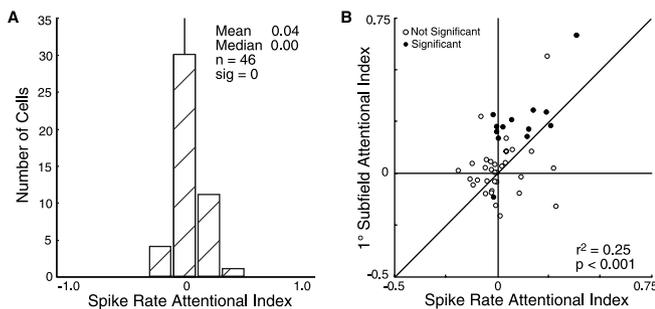


Figure 8. *A*, Histograms of the attentional modulation index for the overall firing rate in response to the stimulus. The attentional modulation index was defined as the spike rate in the attended mode minus the spike rate in the unattended mode divided by their mean. No cells showed individually significant changes in spike rate. *B*, The attentional index for the primary subfield is plotted against the attentional index for the overall spike rate in response to all stimuli. The 13 cells with significant effects on the primary subfield are shown as filled circles.

tion below threshold results simply in the absence of spikes. To a first approximation, however, the firing of the simple cell can be considered to be a linear system followed by an output nonlinearity, as has been suggested for simple cells in the cat (Carandini and Ferster, 2000). The output function translates the result of the linear stage into spikes. If one assumes such a model, the shape of the output nonlinearity can be derived in a straightforward manner (Hunter and Korenberg, 1986). First, the result of the linear stage can be estimated by correlating the stimulus with the receptive field, the linear first-order kernel (Fig. 9). This calculation yields a time-varying linear response that is fed into the output nonlinearity. The shape of the output nonlinearity can be estimated by comparing the actual spiking behavior with the estimated linear response of the first stage (Chichilnisky, 2001). These output nonlinearity curves (Fig. 9) are similar to contrast-response functions in that they measure gradations of the responses of the neuron to systematically varied stimuli. Our stimuli did not vary in contrast, however, but in how well each spatiotemporal arrangement of pixels matched the linear estimation of the receptive field of the neuron. The response thresholds of these curves therefore measure the point at which specific stimulus configurations consistently evoked an appreciable modulation of the firing by the neuron.

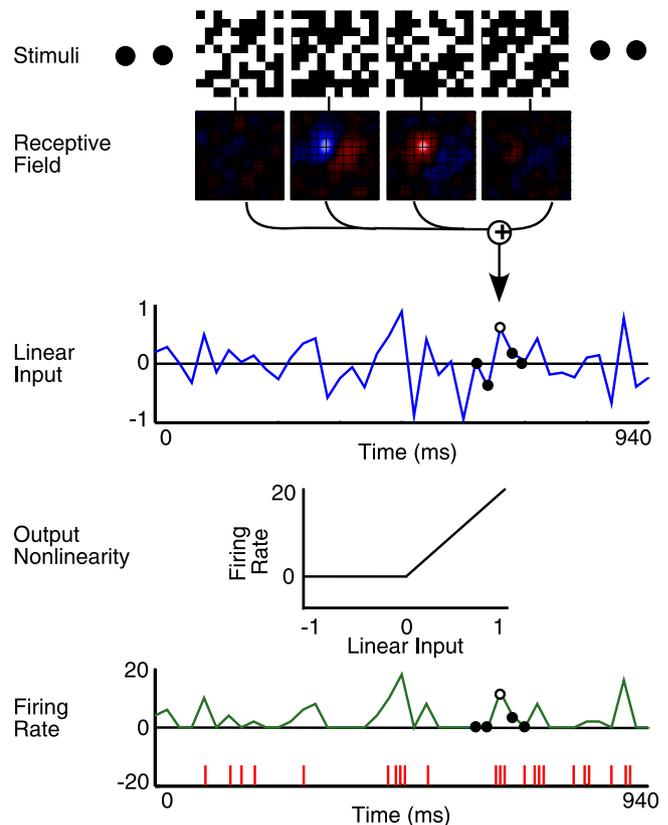


Figure 9. Computation of the two-stage model: a linear filter followed by a nonlinear output function. First row, Schematic representation of four successive stimulus frames, with additional frames indicated by black dots. Second row, The receptive field of a single neuron is convolved with the stimulus sequence to provide a single numerical value representing the result of the linear first stage (the “linear input”). Third row, A portion of the calculated output of the linear first stage, arbitrarily scaled from -1 to 1 ; the x -axis is time and covers 40 different stimulus frames at 23.5 ms/frame. The open circle is the value derived from the stimuli illustrated in the first row; the neighboring black dots indicate the values for the adjacent stimulus frame sequences. Fourth row, A simple form of output nonlinearity: a rectifier that converts all negative inputs to zero and linearly scales positive inputs. Fifth row, Green, The predicted firing rate, derived by passing the result of the linear stage through the output function. The lines shown in red, on the bottom axis, are the actual spikes of the neuron, which varied from zero to three spikes per frame within this specific sequence and up to five spikes per frame overall. The neuron fired more spikes when the predicted firing rate function is high and few or no spikes when the predicted firing rate is low.

When we performed this two-stage analysis of the simple-cell responses, we found that the response nonlinearities were highly variable. For most of our neurons (as in Fig. 10*A*), the nonlinearities resembled a half-wave rectifier, like that drawn in Figure 9. For negative inputs, the average spike rate was near zero; for positive inputs, the output spike rate was an increasing function. In other cases, typically when the neurons had very high firing rates in response to the white-noise stimulus, the output functions were smoothly increasing functions of their linear inputs. It is important to recognize that a wide range of firing rates was elicited in response to our white-noise stimulus. Often, no spikes (32 of 46 neurons) were fired in response to the worst stimuli. However, when the best linear inputs were provided, the white-noise stimulus drove all the neurons vigorously. Across all neurons, the mean firing rate in response to stimuli in the top linear-input bin was 151 Hz (range, 75–425 Hz). Thus, at least some configurations of our white-noise stimuli drove all neurons very strongly.

We asked whether the gain change caused by attention was the

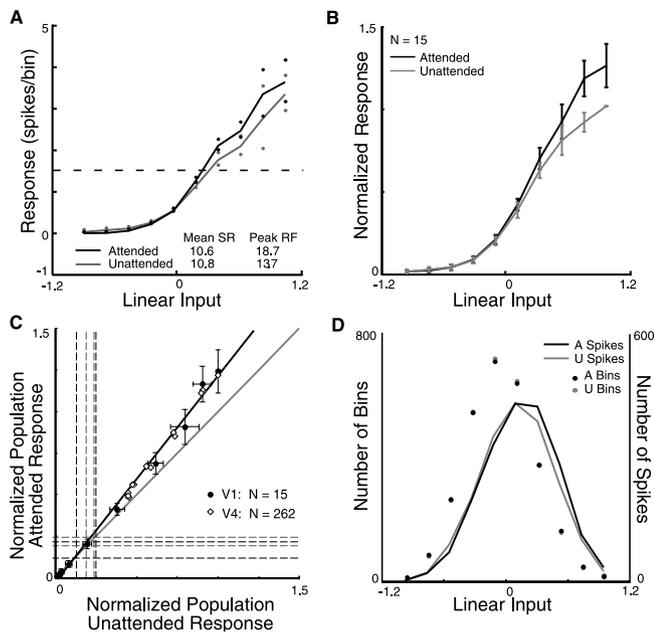


Figure 10. The output functions for a single neuron (**A**) and for the 15 cells with strongest attentional modulation (**B**). The output functions (black, attended mode; gray, unattended mode) represent the firing rate (y -axis) averaged over all frames for which the linear first stage took on a particular range of values (x -axis). In **A**, the dots indicate the SEM for each bin. The dashed line is the average undriven response of the neuron during the fixation period (no stimuli in receptive field). In **B**, the linear inputs from each neuron were aligned on the maximum value of its linear first stage, and the axis was fixed from -1.0 to 1.0 . The responses of each neuron were normalized to the peak response in the unattended mode before averaging across the neurons. The error bars indicate the SE in the normalized bins across neurons. In **C**, the normalized attended response is plotted against the normalized unattended response for the 15 cells with the strongest attentional modulation (filled circles). The black line shown is the linear regression of the attended responses on the unattended responses ($y = 1.24x - 0.02$; $r^2 = 0.994$). The gray line ($y = x$) is the expected result if there was no effect of attention. The SE bars are plotted at each point, for both the attended and unattended response values. The dashed black lines show ± 1 SEM of the normalized undriven activity. Renormalized data from area V4 from McAdams and Maunsell (1999) has been plotted using open diamonds for the response values and dashed gray lines to indicate the undriven activity in that study (see Results). In **D**, we show the population average of the distributions of the number of times bins (circles) and total spikes (lines) occurred for each value of the linear input for the same 15 cells as in **B** and **C** (black, attended; gray, unattended). There were many more time bins, and subsequently more total spikes, for linear input values near zero; the most effective stimuli occurred only rarely and evoked relatively few total spikes.

result of a decreased threshold at which the stimulus caused a neuronal response or from an increase in the slope of the output function. Although our estimates of the output functions were too noisy to make such a distinction for single neurons (Fig. 10A), the population averages for the attended and unattended modes can provide some insight. We restricted our analysis to the 15 neurons that showed the strongest attentional effects on the visual responses (Fig. 10B). For these neurons, the slope of the averaged output nonlinearity changed without changing the threshold. This indicates that, when strongly excitatory stimuli were presented (the top four values of the linear input), the firing rate of the neuron was increased by attention, but the responses to less effective stimuli were changed relatively little. A similar result was seen for the entire population, but the lesser effects of attention in the overall population make the slope change less salient.

There is an apparent conflict between these observed responses to effective stimuli and the minimal spike rate changes we

observed overall. This can be reconciled by considering the distribution of the total number of spikes rather than the spike rates. In Figure 10D, we show the distributions of the number of times bins (circles) and total spikes (lines) that occurred for each value of the linear input in the attended and unattended modes, averaged across the same 15 cells. The output nonlinearity (Fig. 10B) is derived from the ratio of these values (total spikes/number of bins). Because the linear input is the weighted sum of a random stimulus consisting of $+1$ and -1 s, it approximates a Gaussian distribution centered on zero. Although small values of the linear input drove the cells weakly, they were far more common than the most effective stimuli, and thus they evoked a larger total number of spikes. Because attentional effects (Fig. 10B) were more pronounced for the more effective stimuli, attentional modulation of the average spike rate was diluted by the preponderance of spikes evoked by weaker stimuli (Fig. 10D).

Because the weakly driven responses showed little or no attentional modulation, we divided the spikes during visual stimulation into two categories based on whether the average response rate was under or over the firing rate in the absence of visual stimulation. For example, in Figure 10A, the dashed horizontal line is the average undriven activity. The “under” responses were computed by averaging the spikes from the linear-input bins containing responses below this line (for this neuron, the responses to stimuli in the first six bins), and the “over” responses were computed by averaging the spikes from the linear-input bins containing responses above this line (for this neuron, the responses to stimuli in the last four bins). We found that the visually evoked neuronal activity above the undriven firing rate was significantly changed by attention ($n = 46$; paired t test, $p = 0.01$; unattended mean response, 59.3 spikes/s; attended mean response, 62.7 spikes/s). Of the 46 neurons we recorded, five always responded more to the white-noise stimuli than to the background, and seven never spiked at all in response to both the background and some linear-input levels. Excluding these neurons, the visually evoked neuronal activity below the undriven firing rate was not significantly changed by attention ($n = 34$; paired t test, $p = 1.0$; unattended mean response below undriven activity, 4.5 spikes/s; attended mean response below undriven activity, 4.1 spikes/s).

If attention acts multiplicatively on the neuronal responses, a plot of the average responses in the attended mode versus the unattended mode should be well fit by a straight line (Fig. 10C, filled circles). The slope of the linear regression of the attended responses on the unattended responses is 1.24, with a correlation coefficient of 0.994. This indicates that attention increased the visually driven responses in the 15 cells by an average of 24%. Note that the regression line is dominated by the top five values of the spike rate, which were evoked by the highest values of the linear input. These were the only responses that exceeded the average response of these 15 cells in the absence of a visual stimulus (Fig. 10C, dashed lines). When the same analysis was performed on the entire population, the slope of the linear regression was 1.11, with a correlation coefficient of 0.998. The high correlation coefficients of these regressions support the hypothesis that attention causes a multiplicative scaling of the visually evoked neuronal responses.

The multiplicative scaling of orientation tuning curves in area V4 was described previously by McAdams and Maunsell (1999). To allow direct comparison of their results with these, we have renormalized their attended and unattended responses to different orientations so that the response to the peak orientation in the unattended mode is set to 1.0 and then plotted their data as open diamonds, and the undriven activity in that study is illustrated

with gray dashed lines in Figure 10C. Clearly, the points from both studies fall on nearly the same line. The slope of a line fitted to the V4 data was slightly greater ($y = 1.32x - 0.09$; $r^2 = 0.996$), indicating more attentional enhancement in that study. Notably, none of the orientations sampled in the previous study resulted in a mean firing rate below the undriven activity of the neurons, whereas in this study, almost half of the stimuli resulted in inhibition of the neuronal responses.

To further examine the hypothesis of multiplicative scaling for single neurons, we asked how much of the variance in the visual responses in the unattended and attended modes could be accounted for by multiplicative scaling. We scaled the unattended response in the primary subfield by the ratio of the total visual responses in the attended and unattended modes (see Materials and Methods) to yield a predicted attended response. For instance, we scaled the blue curve in Figure 2C to make an approximation to the red curve. We then calculated the variance between this predicted attended response and the measured attended response (using only the bins with a significant response). This was compared with the variance between the measured unattended and attended responses. For the 13 neurons with significant effects of attention on their visual responses, a single scale factor accounted for a large amount of the variance between the measured curves (average, 59%; median, 62%; range, 21–75%).

Discussion

Attention in primary visual cortex

Although attentional modulation has been observed previously in primary visual cortex using functional magnetic resonance imaging (MRI) (Tootell et al., 1998; Brefczynski and DeYoe, 1999; Gandhi et al., 1999; Huk and Heeger, 2000) and electrophysiology (Motter, 1993; Roelfsema et al., 1998; Vidyasagar, 1998; Ito and Gilbert, 1999; McAdams and Maunsell, 1999), the origins of attentional modulation remain elusive. Feedback from extrastriate areas has frequently been postulated as the mechanism for attentional modulation of V1 neurons (Roelfsema et al., 1998; Vidyasagar, 1998; Ito and Gilbert, 1999). Here, we have shown that attentional modulation reaches to an early stage of visual processing in V1: the visual response of simple cells.

The responses of a simple cell may derive from either afferent inputs from the lateral geniculate nucleus (LGN) or from intracortical interactions, including feedback. Attentional modulation of the human LGN has been shown using functional MRI (O'Connor et al., 2002). This suggests the possibility that attention may enhance feedforward sensory input to simple cells from either the LGN or cells in layer 4 of area V1. However, in our experiment, a time-varying stimulus was continuously presented to striate neurons. Therefore, cortical neurons at later levels of processing were activated by the stimulus throughout the periods used to create the simple cell impulse response functions. Feedback from extrastriate cortical inputs could be the source of the observed attentional modulation. Cortical inputs to simple cells have been shown to modulate the visual responses generated by afferent thalamic inputs without altering the orientation selectivity (Ferster et al., 1996; Chung and Ferster, 1998) or the receptive-field structure of simple cells (Kara et al., 2002). Our results are consistent with these types of actions of cortical inputs on simple cells: attention altered the magnitude but not the structure of the receptive field of the simple cell.

Spatial attention

On a trial-by-trial basis, the animal's task was to detect and discriminate the color of a patch appearing in only one of two stim-

uli. In this way, attention was directed to both a particular spatial region, one of the white-noise stimuli, and a particular feature, color. The white-noise stimuli also carried multiple roles in our experiment. The primary role was to map and measure the visual responses. Secondly, they provided the cue to direct spatial attention to a particular region. Finally, they acted as a mask that increased the difficulty of identifying the color of the patch. If the effects of spatial attention predominated our measured responses, we would anticipate an enhancement of the visual response, such as we observed. However, if feature-directed attention were predominant in the simple cell response, we might have seen an inhibition of the visual response because the achromatic mapping stimuli acted as mask with regard to the color-detection task. Ideally, the responses to the color patches might have provided information about attentional changes in the color-detection ability of the simple cells. However, to ensure that spatial attention was distributed across the entire white-noise stimulus, the colored patch was randomly assigned to any pixel within that stimulus and only rarely appeared in the receptive field. Therefore, the data can only demonstrate an effect of spatial attention on simple cell responses in V1 and cannot rule out potential roles for feature-directed attention. More pronounced enhancements of firing rate of V1 neurons, although not specifically simple cells, have been observed in tasks in which the effects of spatial and feature-directed attention might potentiate rather than interfere (Motter, 1993; Roelfsema et al., 1998; Vidyasagar, 1998; Ito and Gilbert, 1999; McAdams and Maunsell, 1999).

Our task required the animal to distribute his attention across the entire stimulus region, which was always larger than the receptive field. Therefore, we would not expect, and did not see, decreases in the receptive-field size with attention. Attention appeared to produce similar effects on each of the antagonistic subfields, suggesting that it does not differentially affect the mechanisms or pathways that create the underlying spatial structure of the receptive field of the simple cell. Both subfields are derived from positive spiking events in reverse correlation analysis: the "on" subfield is the spatial receptive field region in which spikes are initially correlated with white stimuli, and the "off" subfield is the region in which the spikes are initially correlated with black stimuli. Thus, the enhancement of both subfields by attention is not paradoxical and implies a net increase in stimulus-evoked spikes. These results are consistent with the theory that attentional modulation can cause a multiplicative scaling of the response of a neuron to a visual stimulus, as demonstrated previously for orientation tuning in area V4 (McAdams and Maunsell, 1999) and for direction tuning in the middle temporal area (Treue and Martinez Trujillo, 1999; Recanzone and Wurtz, 2000), and further applies multiplicative scaling to the actions of spatial attention on the subregions within a receptive field.

Selectivity of attentional modulation

Most previous electrophysiological studies of attention in primary visual cortex have reached conclusions about attention based on the strength of firing to simple stimuli (Motter, 1993; Luck et al., 1997; Roelfsema et al., 1998; Vidyasagar, 1998; Ito and Gilbert, 1999; McAdams and Maunsell, 1999; Marcus and Van Essen, 2002) rather than examining the detailed coherence between the stimulus and the neural response. Consistent with many of these studies (Motter, 1993; Luck et al., 1997; McAdams and Maunsell, 1999; Marcus and Van Essen, 2002), we found little change in the overall firing rates of the neurons. Aside from a few rare, highly effective stimulus configurations, most frames of the white-noise stimulus that we used excited or inhibited the

neurons weakly (Fig. 10D). Insofar as attention may decrease responses to ineffective stimuli (Fig. 10A, bins 1–3) as well as increase responses to the rarer effective stimuli, our stimulus could result in no change in overall firing rate. However, the attentional effects on the output functions help to clarify this result: attention appears to have much more effect on responses to stronger stimuli than on the responses that were near or below threshold.

In past studies, the most robust attentional effects in striate cortex have been found using stimuli likely to evoke stronger responses in V1 neurons. Many visual neurons have strong surround inhibition, and stimuli close in size to the classical receptive field of the neuron typically elicit the larger responses. Consistent with this, lines close in size to a V1 receptive field have generated the largest reported effects (Motter, 1993; Roelfsema et al., 1998; Ito and Gilbert, 1999), whereas studies using stimuli larger than a V1 receptive field have generally observed smaller attentional effects (Luck et al., 1997; McAdams and Maunsell, 1999; Marcus and Van Essen, 2002). Our stimuli varied greatly in their efficacy; some frames of the white-noise stimulus matched the receptive field well and therefore evoked strong responses, but most frames did not. We found that the better a specific stimulus matched the receptive field, the more apparent the effects of attention (Fig. 10A,B). It should be emphasized that the most effective configurations in the white-noise stimuli drove neurons robustly (average, 151 Hz; range, 75–425 Hz).

The average firing rate of a neuron in response to a particular stimulus is inadequate for assessing how well a neuron is driven because the dynamic range of the responses of that neuron remains unknown. Specifically, a single response does not tell us where it lies relative to the thresholds of the neuron for that type of stimulus and its output nonlinearity function. Although we can never know the absolute upper and lower limits of the responses of a neuron when sampling with a limited number of stimuli, by measuring its responses to multiple stimuli, its output nonlinearity function and threshold can be estimated. For the white-noise stimuli in our experiment, most simple cells had a high threshold for response, generally increasing their firing rate over their undriven rate for only a small percentage of the stimulus configurations. Interestingly, we also found that effects of attention were restricted to only those stimuli. This provides insights into two aspects of the interaction between sensory inputs and their cognitive modulation. From a mechanistic perspective, it suggests that attentional inputs are modulating the visual inputs rather than creating spikes independent of the visual inputs. Attention does not change the threshold of the simple cell for responding but enhances those visual responses that have already surpassed the threshold. Second, from a computational perspective, the responses of a single neuron that are below threshold are most likely ignored (as the attentional inputs are ignored) in those calculations that lead us to a cohesive internal representation of the visual world. Numerous psychophysical experiments have demonstrated the existence of inattention blindness, an inability of observers to perceive unexpected objects in plain view when attention is directed elsewhere (Mack et al., 1992; Simons and Chabris, 1999; Most et al., 2005). These profound effects of attention on consciousness suggest an absence of stimulus information at some level in the visual processing hierarchy. Our results indicate that this selective interpretation of the visual environment may develop from the behavioral gating of visual stimuli at a very early stage of cortical processing, simple cells in V1.

References

- Brefczynski JA, DeYoe EA (1999) A physiological correlate of the “spotlight” of visual attention. *Nat Neurosci* 2:370–374.
- Carandini M, Ferster D (2000) Membrane potential and firing rate in cat primary visual cortex. *J Neurosci* 20:470–484.
- Cave KR, Bichot NP (1999) Visuospatial attention: beyond a spotlight model. *Psychon Bull Rev* 6:204–223.
- Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. *Network* 12:199–213.
- Chung S, Ferster D (1998) Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. *Neuron* 20:1177–1189.
- Connor CE, Gallant JL, Preddie DC, Van Essen DC (1996) Responses in area V4 depend on the spatial relationship between stimulus and attention. *J Neurophysiol* 75:1306–1308.
- Connor CE, Preddie DC, Gallant JL, Van Essen DC (1997) Spatial attention effects in macaque area V4. *J Neurosci* 17:3201–3214.
- Cook EP, Maunsell JHR (2004) Attentional modulation of motion integration of individual neurons in the middle temporal visual area. *J Neurosci* 24:7964–7977.
- Dow BM, Snyder AZ, Vautin RG, Bauer R (1981) Magnification factor and receptive field size in foveal striate cortex of the monkey. *Exp Brain Res* 44:213–228.
- Ferster D, Chung S, Wheat H (1996) Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* 380:249–252.
- Gandhi SP, Heeger DJ, Boynton GM (1999) Spatial attention affects brain activity in human primary visual cortex. *Proc Natl Acad Sci USA* 96:3314–3319.
- Grunewald A, Bradley DC, Andersen RA (2002) Neural correlates of structure-from-motion perception in macaque V1 and MT. *J Neurosci* 22:6195–6207.
- Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol (Lond)* 160:106–154.
- Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol (Lond)* 195:215–243.
- Hubel DH, Wiesel TN (1974) Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor. *J Comp Neurol* 158:295–305.
- Huk A, Heeger DJ (2000) Task-related modulation of visual cortex. *J Neurophysiol* 83:3525–3536.
- Humphreys G, Bruce V (1989) Visual cognition. London: Erlbaum.
- Hunter IW, Korenberg MJ (1986) The identification of nonlinear biological systems: Wiener and Hammerstein cascade models. *Biol Cybern* 55:135–144.
- Hurlbert A, Poggio T (1985) Spotlight on attention. *Trends Neurosci* 8:309–311.
- Ito M, Gilbert CD (1999) Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron* 22:593–604.
- Kara P, Pezaris JS, Yurgenson S, Reid RC (2002) The spatial receptive field of thalamic inputs to single cortical simple cells revealed by the interaction of visual and electrical stimulation. *Proc Natl Acad Sci USA* 99:16261–16266.
- Luck SJ, Chelazzi L, Hillyard SA, Desimone R (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J Neurophysiol* 77:24–42.
- Mack A, Tang B, Tuma R, Kahn S, Rock I (1992) Perceptual organization and attention. *Cognit Psychol* 24:475–501.
- Marcus DS, Van Essen DC (2002) Scene segmentation and attention in primate cortical areas V1 and V2. *J Neurophysiol* 88:2648–2658.
- McAdams CJ, Maunsell JH (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441.
- Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. *Science* 229:782–784.
- Most SB, Scholl BJ, Clifford ER, Simons DJ (2005) What you see is what you set: sustained inattention blindness and the capture of awareness. *Psychol Rev* 112:217–242.
- Motter BC (1993) Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J Neurophysiol* 70:909–919.

- O'Connor DH, Fukui MM, Pinsk MA, Kastner S (2002) Attention modulates responses in the human lateral geniculate nucleus. *Nat Neurosci* 5:1203–1209.
- Posner MI, Snyder CRR, Davidson BJ (1980) Attention and the detection of signals. *J Exp Psychol* 109:160–174.
- Recanzone GH, Wurtz RH (2000) Effects of attention on MT and MST neuronal activity during pursuit initiation. *J Neurophysiol* 83:777–790.
- Reid RC, Victor JD, Shapley RM (1997) The use of m-sequences in the analysis of visual neurons: linear receptive field properties. *Vis Neurosci* 14:1015–1027.
- Roelfsema PR, Lamme VA, Spekreijse H (1998) Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395:376–381.
- Schiller PH, Finlay BL, Volman SF (1976) Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J Neurophysiol* 39:1288–1319.
- Simons DJ, Chabris CF (1999) Gorillas in our midst: sustained inattentive blindness for dynamic events. *Perception* 28:1059–1074.
- Tootell RBH, Hadjikhani N, Hall KE, Marrett S, Vanduffel W, Vaughan JT, Dale AM (1998) The retinotopy of visual spatial attention. *Neuron* 21:1409–1422.
- Treue S (2001) Neural correlates of attention in primate visual cortex. *Trends Neurosci* 24:295–300.
- Treue S, Martinez Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579.
- Treue S, Maunsell JH (1999) Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *J Neurosci* 19:7591–7602.
- Van Essen DC, Newsome WT, Maunsell JHR (1984) The visual representation in striate cortex of the macaque monkey: asymmetries, anisotropies and individual variability. *Vision Res* 24:429–448.
- Victor JD (1992) Nonlinear systems analysis in vision: overview of kernel methods. In: *Nonlinear vision: determination of neural receptive fields, function, and networks* (Nabet B, Pinter R, eds), pp 1–37. Cleveland: CRC.
- Vidyasagar TR (1998) Gating of neuronal responses in macaque primary visual cortex by an attentional spotlight. *NeuroReport* 9:1947–1952.