Behavorial/Systems/Cognitive

Orientation-Specific Adaptation in Human Visual Cortex

Geoffrey M. Boynton and Eva M. Finney

The Salk Institute, La Jolla, California 92037-1099

Nearly all methods for analyzing and interpreting functional magnetic resonance imaging (fMRI) data assume that the fMRI signal behaves in an approximately linear manner. However, it has been shown that the mean fMRI response to a pair of briefly presented visual stimuli is significantly smaller than would be expected from the response to a single stimulus. This smaller response could be the result of either a nonlinearity in the fMRI signal or neuronal adaptation. We tested the neuronal adaptation hypothesis by measuring the fMRI response to sequential pairs of sinusoidal gratings that had either the same or orthogonal orientation. The adaptation hypothesis predicts that brain areas with orientation-selective neurons should show a more linear response when the stimulus pair is orthogonal than when the pair is identical. Our results show no orientation-specific adaptation effects in primary visual cortex (V1) but increasing effects along the hierarchy of visual areas (V2, V3, and V4V). A psychophysical contrast detection experiment, using similar oriented gratings as adapters, shows evidence of orientation-specific adaptation in the visual system. These results have implications for the interpretation of rapid event-related fMRI experiments, as well as for recently developed methods that use adaptation as a tool to measure the response properties of underlying neuronal subpopulations.

Key words: fMRI; adaptation; orientation; visual cortex; psychophysics; nonlinearity

Introduction

Despite the widespread use of functional magnetic resonance imaging (fMRI), relatively little is understood about the source of the fMRI signal (for review, see Heeger and Ress, 2002). Previous studies have shown that, to a first approximation, the fMRI response is a linear transformation in time of the underlying neuronal signal (Boynton et al., 1996; Glover, 1999). This fortunate result allows for a simple interpretation of rapid event-related fMRI experiments in which neuronal events are spaced close enough in time that their associated hemodynamic responses overlap. Given a linear system, it can be assumed that the fMRI responses simply add together, so selective subtraction of the response to subsequent events can provide an estimate of the response to a given event (Dale and Buckner, 1997). However, recent studies show that the fMRI response to the second of a pair of briefly presented visual stimuli is significantly smaller than would be expected from the response to a single stimulus alone (Dale and Buckner, 1997; Huettel and McCarthy, 2000). The fMRI response to the second stimulus alone can be estimated by subtracting the response to a single stimulus from the response to the pair. Huettel et al. (2000) show that the estimated amplitude of the response to a second stimulus can be as little as 50% of the first and can be delayed by as much as 6 sec.

There is a neuronal explanation for this apparent nonlinearity that could save the linear fMRI hypothesis. It is well known from nonhuman primate electrophysiological experiments that the neuronal response to a visual stimulus adapts over time (Albrecht et al., 1984; Maddess et al., 1988; Bonds, 1991; Carandini et al., 1998). It is therefore possible that the smaller fMRI response to the second of the paired stimuli is simply attributable to a smaller neuronal response.

We tested the neuronal adaptation hypothesis using an event-related fMRI paradigm in which we compared the fMRI response to pairs of sinusoidal gratings that had either the same or orthogonal orientations. We found a strong suppression and delay in the response to the second stimulus. This suppression was equal in magnitude for both same and orthogonally oriented grating pairs in V1 and V2 but became increasingly orientation specific from V3 to V4V.

To establish that the brief stimuli used in our fMRI paradigm actually does produce orientation-specific adaptation in the visual system, we measured contrast detection thresholds after subjects viewed grating stimuli identical to that used in the fMRI experiment. This psychophysical experiment showed clear evidence of orientation-specific adaptation in the human visual system.

Our results suggest that at least some of the nonlinearity seen in the fMRI response in higher visual areas can be attributed to neuronal adaptation effects. The remaining nonlinearity in the response found after our manipulation of stimulus orientation may have either neuronal or vascular origins, or both.

Materials and Methods

Subjects. The same four paid subjects [two male, two female (one was an author, E.M.F.), all right-handed, mean age of 30 years] participated in both the fMRI and psychophysical studies. All subjects had normal or corrected-to-normal visual acuity. All subjects indicated informed written consent in accordance with the Salk Institute Human Subjects Review Board guidelines.

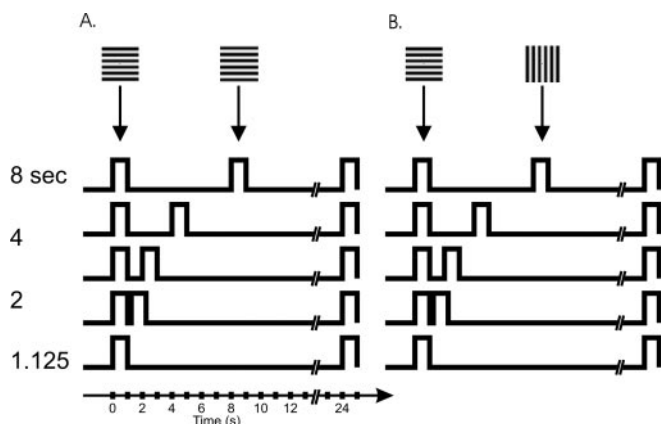*fMRI stimulus presentation apparatus.* Stimuli for the fMRI experi-

Figure 1. Stimulus diagram for fMRI experiments. Subjects viewed 1 sec duration counter-phase-modulated gratings either alone or in successive pairs. Stimulus onset asynchronies were (from top to bottom) as follows: 8, 4, 2, and 1.125 sec. A new trial began every 24 sec. The first stimulus could be either vertical or horizontal (data not shown). Pairs had either the same orientation (*A*) or were oriented orthogonally (*B*).

ments were generated on an Apple PowerMac G3 laptop computer (Apple Computers, Cupertino, CA) using Matlab version 5.2 (MathWorks, Natick, MA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Images were produced using a Proxima DP9300 LCD projector (In Focus, Wilsonville, OR) fitted with a zoom lens and projected onto a back projection screen. Viewing distance was 22 cm. Projector refresh rate was 60 Hz, and the mean luminance was 100 cd/m^2. Projector gray scale values were linearized with respect to luminance using a photometer.

During fMRI data acquisition, subjects viewed the image on a screen near the subject's neck through a mirror mounted to the MRI table above the subject's eyes. A bite bar stabilized the subject's head.

fMRI stimuli. All stimuli used in the fMRI experiments were 8 Hz counter-phase-modulated sinusoidal gratings of 100% contrast. Gratings were oriented either vertically or horizontally, with a spatial frequency of 0.25 cycles/°, subtending a visual angle of 30° and lasting 1 sec. This low spatial frequency was used because limitations in the presentation apparatus reduced the contrast of higher spatial frequencies. A square fixation point was placed in the center of the visual field, and a 2° gray border with the same mean luminance of the grating surrounded the stimulus. Trials in the fMRI experiment consisted of either a single grating stimulus or pairs of grating stimuli separated with a stimulus onset asynchrony (SOA) of 1.125, 2, 4, or 8 sec (Fig. 1). Each stimulus in the pair was oriented either vertically or horizontally so that the pair had either the same (Fig. 1*A*) or orthogonal (Fig. 1*B*) orientations. A uniform gray field of same mean luminance as the gratings with a central square fixation point was presented between stimulus presentations. Each trial lasted 28 sec. An fMRI stimulus scan consisted of a 10 sec blank period, followed by six trials, lasting a total of 6*28 + 10 = 178 sec.

Each experimental session consisted of a reference scan (see below), followed by 12 stimulus scans. The SOA was held constant within an experimental session. Six stimulus types were shown within each scan, made up of two adapting orientations crossed with three stimulus conditions (adapter alone, same, and orthogonal pairs). The order of the six stimulus types was counterbalanced across the 12 stimulus scans.

In a subset of the fMRI scans, subjects performed a task at fixation to maintain spatial attention, vigilance, and eye fixation. The task resembled the memory game "Simon" that involved memorization and recall of spatiotemporal patterns. During these experimental runs, subjects responded using a 10 key response box that was electronically isolated from the laptop computer to minimize radiofrequency interference. The task began with a manipulation of the central 3 × 3 pixels in the image (0.5°), which covered the same area as the fixation point in the other experiments. Within this 3 × 3 array of pixels, a pixel would flash from gray to white for 1 sec, followed by another randomly selected pixel, beginning with a three-element sequence for the subject to memorize. After this display period, subjects attempted to repeat the sequence on a

3 × 3 array of keys on a numeric keypad. Feedback was given after each key press by flashing the appropriate pixel within the 3 × 3 array at fixation. At the end of each recalling period, the sequence would be presented again with one more element appended to the end. If an error occurred during recall, the task would reset to a new three-element sequence. Because the subject controlled the task, the timing was uncorrelated with the presentation times of the full-field test stimuli. No significant difference was found in the fMRI responses to the test stimuli between the Simon task and passive viewing conditions.

An alternate method of controlling for attention would be to direct attention toward the test stimuli instead of away toward fixation. This would have the advantage of producing a larger response, attributable to the added attentional effect. However, because the test stimulus was only present for a small proportion of the time, there would be little control over the subjects' attentional state during most of the scanning period.

fMRI data acquisition. Functional MRI data were acquired using a Siemens (Munich, Germany) Vision 1.5 T scanner at Thornton Hospital at the University of California, San Diego using a small flex coil and a low-bandwidth echo-planar imaging sequence. During each scan, 178 temporal frames were acquired over 178 sec (repeition time, 1 sec; flip angle, 50°; eight coronal slices of 6 mm thickness and 4 × 4 mm resolution; field of view, 256 mm). The first 10 temporal frames (10 sec) were discarded to avoid magnetic saturation effects. Each scanning session ended with an anatomical scan (MPRAGE; 1 × 1 × 1 mm resolution), using a standard T1-weighted gradient echo pulse sequence. Anatomical scans were used to align functional data across multiple scanning sessions to a subject's reference volume. The timing of stimulus presentation was synchronized with fMRI data acquisition using a custom-made phototransistor trigger that detected a light indicator on the scanner keyboard.

fMRI region of interest selection. Occipital visual areas V1, V2, V3, and V4V were defined using standard retinotopic mapping and cortical-flattening techniques as described previously (Engel et al., 1994; Sereno et al., 1994; Boynton et al., 1999). Regions of interest within these predefined visual areas were selected by means of a reference scan that was run at the beginning of each session. For the reference scan, a plaid stimulus (superimposed vertical and horizontal gratings) of the same size, spatial frequency, and temporal frequency as the test stimuli was presented in alternation with a uniform gray field for six blocks of 28 sec cycles. For subsequent analysis, we chose voxels that correlated ($r > 0.35$) with a six cycle sinusoid and had a temporal phase lag between 0.6 and 8 sec.

This procedure restricted the analysis to a subset of the original contiguous set of voxels covering each visual area. This restriction ranged from 15 to 35% of total voxels across subjects. The amount of restriction was approximately the same across visual areas, indicating that our reference scan did not bias our analysis toward one visual area over another.

fMRI data analysis. Linear trends were subtracted from the fMRI time course of each voxel and were then divided by the mean to convert our measure of responses to percentage signal change. Time courses of voxels within a given visual area and responses to the reference stimulus were averaged for each stimulus condition across scans and sessions. Vertical and horizontal stimulus conditions were averaged together resulting in three stimulus conditions per SOA: a single stimulus condition, a same condition in which the pair of stimuli had the same orientation, and an orthogonal condition in which the pair had orthogonal orientations. We found no significant difference between the results from our subjects' left and right visual cortices, so signals from the two hemispheres were averaged together.

Smooth curves were fit to the fMRI responses in a given region of interest using a parametric function, which is the difference of two gamma functions:

$$h(t) = h_1(t) - h_2(t) + k,$$

where:

$$h_i(t) = \frac{((t - \delta_i)/\tau_i)^{(n_i-1)} e^{-(t-\delta_i)/\tau_i}}{\tau_i(n_i - 1)!}$$

The parameters n_1 *and* n_2 were set to 4 and 7, respectively, and the remaining five parameters were allowed to vary freely for each curve,

except for the baseline parameter k, which was fixed across all conditions and SOAs. Note that each component of this function, h_1 and h_2, has the same parametric form as used in previous fMRI studies (Boynton et al., 1996). The difference between two gamma functions predicts a biphasic response when the second function, $h_2(t)$, has a slower time course than the first. This smooth parametric function provides a continuous description of the curves and provides an estimate of parameters such as maximum height and delay.

Psychophysical stimulus presentation apparatus. Psychophysical contrast detection experiments were performed outside the scanner in the laboratory under conditions designed to match the fMRI stimulus conditions. Subjects sat upright in a dark room and used a chinrest to maintain the viewing distance of 22 cm from a computer monitor (Multiscan 500Ps; Sony, Tokyo, Japan). Stimuli were generated using an Apple G4 computer equipped with a 10 bit video card (Radius, Sunnyvale, CA) providing 1024 gray scale values. Monitor gray scale values were linearized with respect to luminance using a photometer. As in the fMRI studies, the video refresh rate was 60 Hz, and the mean luminance of the stimulus was 100 cd/m^2.

Psychophysical stimuli. The adapting stimulus in the psychophysical experiments was identical to that of the fMRI experiments (30 × 30° counter-phase-modulated grating; 0.25 cycles/°, 1 sec duration). Unlike the fMRI experiments, test stimuli were Gabor patches, which were 0.25 cycles/° sinusoidal gratings windowed by a Gaussian with a 1/e half-width of 1°. The Gabor patches were centered either 5° above or below the fixation point, falling within the region covered by the adapting stimulus.

Contrast detection thresholds were measured for test stimuli presented either 1.125, 2, 4, or 8 sec after the onset of a 100% contrast 8 Hz counter-phase-modulated grating that had the same spatial frequency, orientation, and size as those used in the fMRI experiments.

Psychophysical methods. Contrast detection thresholds were obtained using a three-down, one-up staircase procedure with a two-alternative spatial forced-choice trial structure. A 500 msec auditory beep preceding each test stimulus by 250 msec alerted the subject to the test stimulus. On each trial, the subject's task was to indicate the location of the test stimulus (above or below fixation) with a key press. Subjects had to respond within a 500 msec response interval. Otherwise, the trial was discarded. The contrast on the next trial was reduced after three correct responses in a row and increased after a single incorrect response. Each staircase consisted of 40 trials. Psychometric functions (percentage correct as a function of contrast) were fit with a Weibull function using a maximum likelihood estimate. The contrast that predicted 80% correct performance was defined as the contrast detection threshold.

For the baseline contrast detection thresholds, which were obtained without an adapting stimulus, new trials began every two seconds. In the adapting conditions, the adapting stimuli were identical to those used in the fMRI experiments. For the SOA of 1.125 sec, no auditory beep was used, and a new adapting stimulus began every 8 sec (Fig. 2A,B). Detection thresholds for the other SOAs were obtained using the same adapting stimulus. That is, after the adapting stimulus, test trials were presented 2, 4, and 8 sec afterward (Fig. 2C,D). A new adapting stimulus began every 16 sec.

Thresholds were measured 12–16 times per condition. Thresholds falling outside 1.5 SDs from the mean were excluded to provide a robust estimate of average threshold. The adaptation ratio for each SOA was defined as the ratio of the contrast detection threshold for that SOA and the baseline contrast detection threshold. Error bars of the adaptation ratio were computed using a bootstrapping procedure.

Results
fMRI results
Figure 3 shows the time course of the fMRI signal to the three conditions in V1 averaged across all four subjects. The diamonds indicate the response to a single stimulus, squares indicate the same-paired stimuli, and circles correspond to the orthogonally paired stimuli. The vertical gray shaded regions denote the time the stimuli were presented. The smooth curves are the best fit of the difference of gamma functions to each curve, as described above.
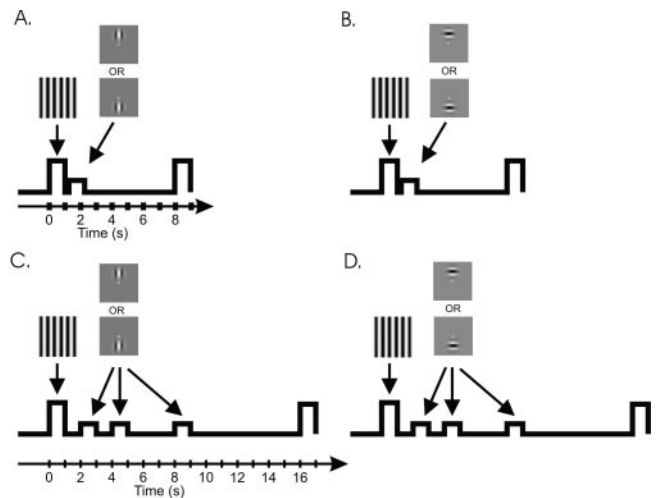

Figure 2. Stimulus diagram for psychophysical experiments. Subjects made a spatial two-alternative forced-choice decision about the location of the target stimulus, which was a small Gabor patch presented either above or below fixation. The test interval containing the Gabor patches followed a 1 sec sinusoidal grating stimulus (identical to that used in the fMRI experiments) with an SOA of 1.125, 2, 4, or 8 sec. Gabors had either the same orientation as the adapting stimulus (A, C) or were orthogonal (B, D). For the SOA of 1.25, a new trail began every 8 sec (A, B). For the other SOAs, trails began every 16 sec (C, D). For these trials, subjects were tested three times for each adapting stimulus, once at each of the SOAs (2, 4, and 8 sec). In all conditions, an auditory beep alerted the subject to the oncoming test stimulus.
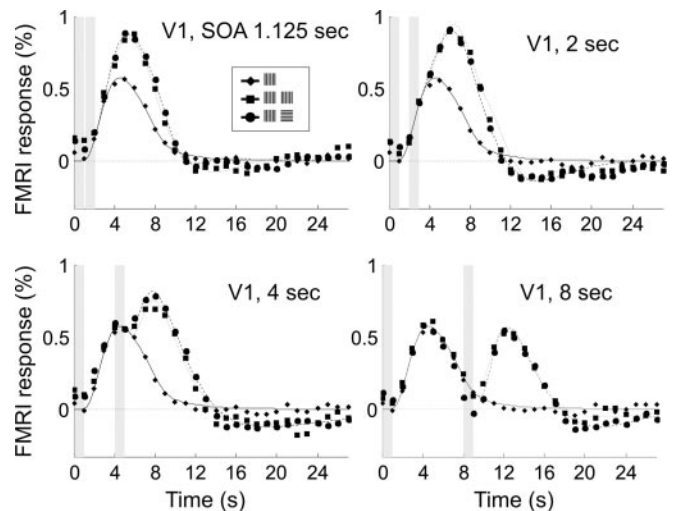

Figure 3. fMRI time courses in V1. Time courses are shown averaged across repetitions and subjects, for the response to a single stimulus alone (diamonds), a pair of same-oriented gratings (squares), and a pair of orthogonal gratings (circles), for each of the SOAs tested (1.125, 2, 4, and 8 sec). Gray vertical bars indicate when the stimuli were presented. Smooth curves through the data are predictions from the best-fitting difference of gamma functions.

As shown by Dale and Buckner (1997) and Huettel et al. (2000), the amount of response suppression to the second stimulus can be seen more easily by subtracting the response to the single stimulus from the other conditions and shifting the difference in response back in time by the SOA. Figure 4 shows the results of this analysis. Here, it is clear that, for short SOAs, the response to the second stimulus is smaller and delayed in time. This difference in amplitude and delay recovers for increasing SOA and has nearly fully recovered when 8 sec separate the onset of the first and second stimulus. Figures 3 and 4 show that, in V1, the relative orientation of the second stimulus has no differential affect on the fMRI response to the second stimulus.
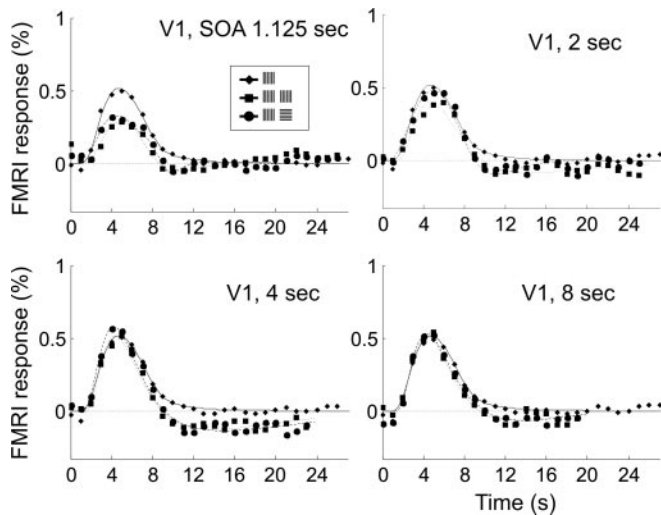
Figure 4. Predicted response to second stimulus in V1. Time courses from Figure 3 after subtracting the response to the single stimulus alone (diamonds) and shifting the response function back in time to account for the SOA. Symbol conventions are as in Figure 3.
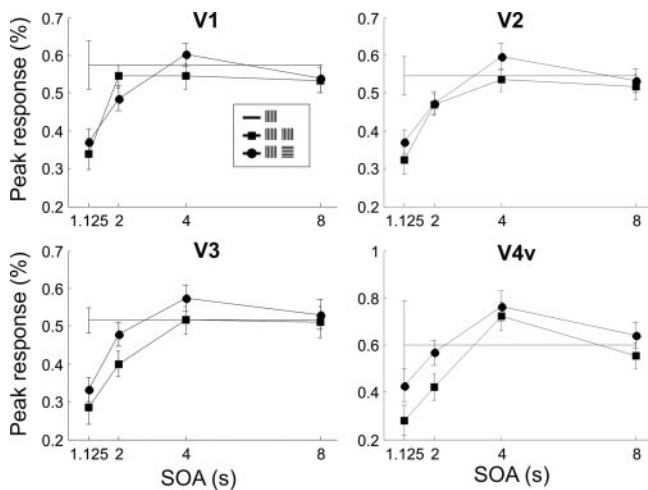

Figure 6. Time-to-peak response (relative to stimulus onset) to the single stimulus alone and to the paired stimuli as a function of SOA in areas V1, V2, V3, and V4V. Symbol conventions are as in Figure 5.

Figure 5. Maximum amplitude of response to the single stimulus alone (horizontal line) and to the response to the second stimulus after subtracting out the first for the same-orientation pair (squares) and orthogonal pair (circles) for responses in areas V1, V2, V3, and V4V as a function of SOA.

Figure 5 shows the time course of the recovery of the response amplitude to the second stimulus across four visual areas (V1, V2, V3, and V4V). These values were calculated by finding the amplitude at the maximum of the best-fitting parametric function described above. Error bars were calculated using a bootstrapping technique in which the original data were repeatedly resampled and refit with the parametric function. This bootstrapping technique provides a method for calculating the SEM for the estimated amplitude. As in Figures 3 and 4, the diamonds in Figure 5 indicate the amplitude of the response to a single stimulus alone, squares indicate the same orientation condition, and circles indicate the orthogonal orientation condition.

In all visual areas, the amplitude of the response to the second stimulus recovers with increasing SOA, so that, after 8 sec, the response to the second stimulus is similar to the response to a single stimulus alone.

Most interestingly, the effect of stimulus orientation varies across visual areas. Whereas there is no significant effect of ori-
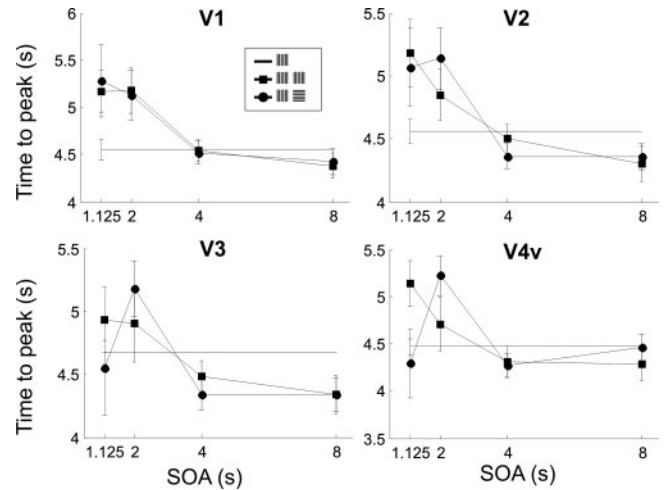
entation on the suppression effect for V1 ($p = 0.46$; repeated-measures ANOVA; main effect for orientation) or V2 ($p = 0.10$), significant orientation-specific adaptation effects are found in areas V3 and V4V ($p < 0.05$ for V3; $p < 0.01$ for V4V).

There were no systematic differences in the results from each of the four individual subjects. Each subject showed a monotonic increase in the orientation-specific effect from areas V1, V2, V3, and V4V, with the exception that two of the subjects showed the same amount of orientation-specific adaptation in V3 and V4V. The time courses of the hemodynamic responses were similar across subjects, so the analysis described above was performed by averaging the hemodynamic responses and then calculating statistics on the average. An alternate analysis was performed in which statistics were calculated for each of the individual subject's hemodynamic responses, and these statistics were then averaged across subjects. This analysis produces a very similar pattern of results; significant orientation-specific adaptation effects were found only in areas V3 and V4V.

Figure 6 shows the time course of the delay of the response to the second stimulus across visual areas. Delays were calculated by finding the time at maximum of the best-fitting parametric functions. Again, error bars were computed using a bootstrapping algorithm. In general, the time-to-peak for the second stimulus is longer than for the single stimulus for shorter SOAs, and this delay decreases with increasing SOAs. This is consistent with the results of Huettel et al. (2000). As with response amplitude, the effect of orientation on the delay of the fMRI response varies across visual areas. In V1, there is no significant difference in the responses to the same versus orthogonal orientation conditions. However, once again, a trend develops across the hierarchy of visual areas. The delay for the shortest SOA of the response to the orthogonal stimulus decreases from V2 to V3 to V4V, until it matches the delay of the response to the single stimulus alone. This results in a non-monotonic delay as a function of SOA for the orthogonally oriented stimulus. In contrast, the response to the same-oriented stimulus recovers in the predicted gradual manner with SOA for all visual areas.

Psychophysical results

Adaptation of the visual system, using stimulus conditions similar to the fMRI experiments, was measured behaviorally using
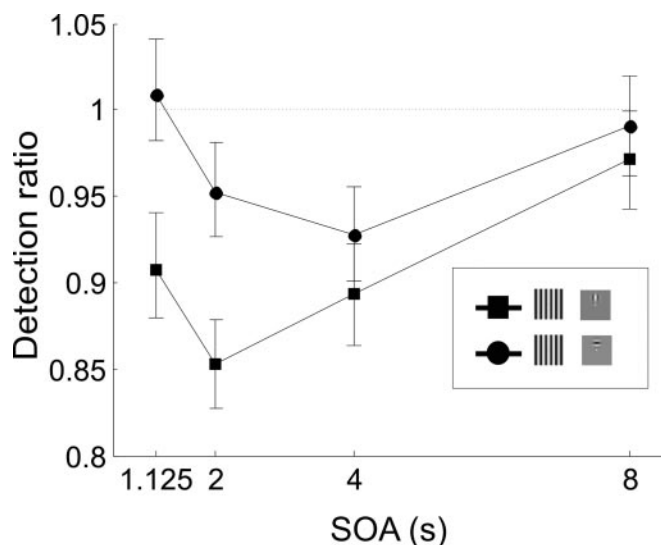
Figure 7. Psychophysical results. Detection ratio is the ratio of the contrast detection threshold without an adapting stimulus to the threshold after an adapting stimulus. Squares indicate adaptation by the same-oriented grating, and circles indicate adaptation after the orthogonal grating.

psychophysical contrast detection thresholds. The minimum contrast needed for detection of a small patch of sinusoidal gratings was measured with and without a preceding full-field adapting grating stimulus, identical to that used in the fMRI experiments. Contrast detection thresholds are a traditional way of probing the adaptive state of the visual system, with the assumption that these thresholds reflect the sensitivity of the population of neurons associated with the detection process (Campbell and Robson, 1968). The ratio of thresholds with and without adaptation is a measure of the fraction that the adapting stimulus reduces the response of the detecting neurons. This measure is therefore analogous to the changes in the fMRI response to the second stimulus relative to the first stimulus shown in Figures 5 and 6.

Figure 7 shows this ratio of thresholds, which we call the detection ratio, averaged across the same four subjects from the fMRI experiments. A detection ratio of 1 implies that the adapting stimulus has no effect on detection thresholds. A detection ratio of ½, for example, suggests that the sensitivity of the neurons responsible for detecting the target have been reduced by a factor of two by the adapting stimulus. Squares are ratios for detecting a grating of the same orientation as the adapting stimulus. Circles are the detection ratios for an orthogonal adapting stimulus. The main result is that detection ratios for the orthogonal targets are consistently higher than for the same-oriented targets; after viewing a 1 sec high-contrast grating, it is easier to see a grating of orthogonal orientation than of the same orientation. The overall effect of SOA on adaptation has an interesting shape; maximum adaptation occurs for both orientations not immediately after the adapting stimulus but ∼2 sec later. After 8 sec, detection ratios for both orientations return to 1.

Discussion

Our psychophysical results demonstrate that the brief stimulus used in the fMRI experiments is capable of adapting the human visual system in an orientation-specific manner. These psychophysical results are consistent with the orientation-specific effects on the fMRI responses in extrastriate cortex. This suggests that at least some of the suppression effect in the fMRI response is the

result of neuronal adaptation and is not attributable to nonlinearities in the neurovascular coupling process. However, why were orientation-specific adaptation effects not found in primary visual cortex? The possible explanations fall into two categories: neuronal and vascular.

One possible neuronal explanation is that a significant proportion of V1 cells have little or no orientation tuning at all (Ringach et al., 2002). Cells in the input layers of V1 (layer 4), receiving signals from the lateral geniculate nucleus (LGN), tend to have broader orientation tuning than cells in the output layers (2, 3, 4B, 5, or 6) that send signals to higher visual areas (Ringach et al., 1997). It is possible that the inclusion of these broadly tuned neurons into the population-based fMRI response may be the cause of the reduced orientation-specific adaptation effect in V1. The similar lack of orientation-selective adaptation in V2 may be because V1 and V2 have similar orientation tuning properties (Levitt et al., 1994).

A second neuronal explanation may be that early visual neurons simply do not adapt in the same way as neurons in higher visual areas. Adaptation effects in the Macaque primary visual cortex are orientation specific, but an orthogonal stimulus nevertheless does produce some adaptation (Carandini et al., 1998). However, very little is known about orientation-specific adaptation properties of neurons outside of V1, mostly because grating stimuli are not often used when recording from extrastriate neurons. So it is possible that, at the single-unit level, orientation-specific adaptation is more effective in extrastriate areas, even if the orientation tuning properties are similar across visual areas.

Finally, the low spatial frequency (0.25 cycles/°) of our stimuli may have favored orientation-selective adaptation in higher visual areas in which receptive fields are larger. Despite this low spatial frequency, our responses are robust across all visual areas. This is consistent with a previous finding that the fMRI response is approximately low-pass with spatial frequency across visual areas (Singh et al., 2000). However, a large fMRI response at 0.25 cycles/° does not mean that the underlying neurons are orientation tuned. Thus, it is possible that, in higher visual areas, there is more selective orientation tuning at 0.25 cycles/°, which would lead to more orientation-specific adaptation than in V1 and V2.

On the vascular side, one reason for our results in V1 may have to do with the source of the blood oxygenation level-dependent (BOLD) signal. A study comparing the BOLD signal with electrophysiological measurements in the macaque visual cortex suggests that the BOLD signal correlates best with local field potentials (LFPs) rather than spiking rates of neurons (Logothetis et al., 2001). LFPs are measured by filtering the electrophysiological signal within the range of ∼40 to 130 Hz. The source of this signal is not well understood, but it is likely driven by rapid synchronized oscillations in membrane potentials and not necessarily by spiking activity. Because it has been hypothesized that LFP signals are driven by input signals into a given area rather than processing within or signals exiting an area, Logothetis and colleagues suggest that the BOLD signal may be the result of presynaptic activity leading into a region of interest rather than spiking activity within an area. The majority of the input to V1 is from the LGN, and LGN cells are known to have little or no orientation selectivity. This could explain why we find very little orientation-specific adaptation in the BOLD signal in V1; LGN signals are not expected to be selective for stimulus orientation. However, it cannot explain why we do not find orientation-specific adaptation effects in V2, because V2 presumably receives orientation-selective inputs from V1.

Another vascular explanation of our V1 results may be that the

blood supply system in the visual cortex is not homogeneous. The input into the primary visual cortex of the squirrel monkey (layer IVc) is more richly vascularized than other layers (Zheng et al., 1991). Thus, the BOLD response may be disproportionately weighted by activity in the input layers, leading to a relatively small amount of orientation-selective adaptation.

Finally, the lack of orientation-specific adaptation in V1 and V2 may be caused by saturation in the hemodynamic response. Suppose the relationship between neuronal and the fMRI response has a saturating nonlinearity. If the response to the stimulus is early visual areas (V1 and V2) is closer to the maximum hemodynamic response than the response in V3 and V4V, then less orientation-specific adaptation might be expected in V1 and V2. This is because an equal size difference in the neuronal response would result in relatively smaller changes in the BOLD signal if the signal were already near saturation. However, the signal changes produced by our brief stimuli, which are on the order of 1%, are substantially smaller than signals produced by longer-duration stimuli in block-design experiments, which can reach 3% with the same imaging equipment.

Delay of the fMRI response with adaptation
Figure 6 shows how the adapting stimulus affects the time-to-peak of the response to the second stimulus. The hemodynamic response can be delayed by almost an additional 1 sec with a short SOA. The neuronal response is not likely to be delayed by this much, so these fMRI results are not consistent with a simple linear convolution of an adapted neuronal signal. For the shortest SOA (1.125 sec), the effect of stimulus orientation on the time-to-peak has the same pattern as that of the amplitudes; responses are increasingly linear from V1 to V4V, only for the orthogonal pair. This suggests that neuronal adaptation at short SOAs not only reduces the amplitude of the BOLD signal but also slows down its time course.

This result has important implications for the interpretation of rapid event-related fMRI experiments. If the amplitude of an event-related signal is calculated by the height of the BOLD signal at a fixed time after stimulus onset, then changes in the time course of the BOLD signal with adaptation may lead to misinterpretation of results.

Adaptation in the fMRI response
Adaptation has been seen in the fMRI response since the earliest use of the technique. In one previous study, Boynton et al. (1996) showed that fMRI responses to short (3 sec) duration stimuli are too large to predict responses to longer stimuli using a linear model. This is consistent with a neuronal response that adapts over time, causing a relatively large neuronal response to short-duration stimuli. Similarly, using a paired-stimulus design, Dale and Buckner (Dale and Buckner, 1997; Buckner, 1998) showed a suppressed response to subsequent stimuli. In a thorough study using the paired-stimulus paradigm, Huettel and McCarthy (2000) found clear adaptation effects that, like the present study, show that the response to the second stimulus is not only smaller but also delayed in time. All of these studies used the same stimulus for adapting as for testing, so they do not address the question of whether these adaptation effects could be occurring within sub-populations of neurons or whether suppression is associated with the neurovascular coupling process.

Pattern specific adaptation and fMRI
There is previous fMRI evidence of pattern-specific adaptation in V1 and higher visual areas. Tootell et al. (1998) measured the fMRI response to the transition between viewing grating stimuli of one orientation to another orientation in V1. They found that the largest transients occur for orthogonal transitions and drop off with decreasing orientation changes. These transients are interpreted as the increase in population response when the presentation of a new oriented stimulus excites a fresh subpopulation of neurons after the previous population has fatigued. The amplitude of this transient as a function of orientation change is interpreted as a reflection of the tuning properties of underlying neurons in V1. However, by design, their responses must reach zero for their zero-orientation transition because there is no stimulus change in this condition. Thus, this study does not reveal the absolute amount of orientation-specific adaptation at each orientation change, but it does show that the relative amount of adaptation is affected by the orientation change.

Unlike the present study, Tootell et al. (1998) did find orientation-specific adaptation in V1. This discrepancy is likely attributable to one or more of a variety of differences between the methods in the two studies. The main difference is the duration of exposure to the adapting stimulus. Whereas our adapting stimulus lasted for 1 sec, the stimulus in the Tootell et al. (1998) study lasted 40 sec. This longer adaptation period is more similar to those used in traditional psychophysical and physiological studies of adaptation. It seems likely that a longer adapter using our paradigm would produce significant orientation-specific adaptation effects in V1. However, it is important to note that, in the present study, significant suppression was found in the fMRI response in V1 for the same-orientation condition, leaving plenty of room for a release from this suppression in the orthogonally oriented condition. Tootell et al. (1998) did not present effects of adaptation in visual areas beyond V1, but, according to our results, their effects should be even larger in V2, V3, and V4V.

Adaptation of the fMRI response can be very robust in higher visual areas and, in fact, has been used as a tool to study the selectivity of neuronal subpopulations (Grill-Spector et al., 1999; Grill-Spector and Malach, 2001). In these studies, the response selectivity of neuronal populations in fusiform brain regions are estimated by measuring the fMRI response to paired stimuli in which the pairs are either identical or differ along some stimulus dimension. For example, in lateral occipital regions known to be responsive to faces, the presentation of a successive pair of identical faces but differing in either size or position produces similar adaptation effects as the same face shown twice (Grill-Spector and Malach, 2001). This is interpreted as evidence that individual units in the underlying brain area are insensitive to the size and position of a face. Interestingly, like the present study, little stimulus-specific adaptation was found in V1, but a significant amount was found in area V4V.

Psychophysical measurements of orientation-specific adaptation
A similar population-based argument has been used to explain pattern-specific adaptation effects measured psychophysically. Viewing a high-contrast grating stimulus elevates the thresholds for detecting subsequently presented gratings, especially if they have similar orientations and spatial frequencies (Blakemore and Nachmias, 1971). The general model is that a test stimulus is detectable if the population response representing that stimulus exceeds a given threshold. Thus, just as with the fMRI studies, adaptation can be used as a tool to probe the pattern specificity of neuronal populations.

Orientation-specific adaptation has been shown to be minimal when the adapter and test are orthogonal (Blakemore and

Nachmias, 1971). For matching adapters and targets, detection thresholds recover as an exponential function over time (Greenlee et al., 1991). The psychophysical results in the present study (Fig. 7) also show orientation-specific adaptation. In our results, the recovery with increasing SOA is non-monotonic; maximum adaptation occurs for intermediate SOAs. It is not clear why these results differ from previous studies. It should be noted that the timing parameters for the shortest SOAs were slightly different than for the longer SOAs.

For longer SOAs, the recovery of detection thresholds shown in Figure 7 follows approximately the same time course as the recovery of fMRI responses to the second stimulus. In particular, the psychophysical results match most closely with the fMRI results in V4V, which show the greatest orientation-specific adaptation effects. An interesting speculation is that the neurons that represent of the stimulus at the time of detection are found in these extrastriate areas.

Implications of the results

Previous paired-stimulus studies of the BOLD signal have been unable to separate vascular from neuronal sources of suppression because the two stimuli in the pair have been identical (Boynton et al., 1996; Dale and Buckner, 1997; Huettel and McCarthy, 2000; Miller et al., 2001). The suppression seen in these previous studies has been generally attributed to hemodynamic factors. On the other hand, a growing number of studies are using paired stimuli to study tuning properties of underlying neuronal subpopulations, a paradigm called fMRI adaptation (Grill-Spector et al., 1999; Grill-Spector and Malach, 2001; Kourtzi and Kanwisher, 2001; Avidan et al., 2002; Epstein et al., 2003). Interestingly, these studies typically attribute suppression effects to neuronal adaptation.

Our results suggest that some of the suppression of the fMRI response to the second stimulus is caused by neuronal adaptation, and some is likely attributable to nonlinearities in the neurovascular coupling process underlying the BOLD signal. The lack of orientation-specific adaptation in early visual areas may mean that the relative contribution of these two effects varies across brain areas. Thus, our results have important implications for the interpretation of all fMRI studies. A nonlinearity in the BOLD signal may have a neuronal sources, and suppression seen in an fMRI adaptation paradigm may have a significant hemodynamic component. Future work is needed to disambiguate these two sources of suppression, which could involve simultaneous electrophysiology and fMRI in nonhuman primates (Logothetis et al., 2001; Logothetis, 2002), additional comparison of fMRI signals and psychophysical results, and a comparison of fMRI and electroencephalogram signals.

Finally, our fMRI and psychophysical results suggest a strategy that early visual areas maintain a veridical representation of a stimulus, whereas higher visual areas represent what is behaviorally relevant. This would include the suppression of responses to repeated stimuli in higher visual areas. That is, our orientation-specific adaptation effects show how visual signals are transformed from low-level sensory responses to higher-level perceptual representations along the hierarchy of visual processing.

References

Albrecht DG, Farrar SB, Hamilton DB (1984) Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. J Physiol (Lond) 347:713–739.

Avidan G, Hasson U, Hendler T, Zohary E, Malach R (2002) Analysis of the neuronal selectivity underlying low fMRI signals. Curr Biol 12:964–972.

Blakemore C, Nachmias J (1971) The orientation specificity of two visual after-effects. J Physiol (Lond) 213:157–174.

Bonds AB (1991) Temporal dynamics of contrast gain in single cells of the cat striate cortex. Vis Neurosci 6:239–255.

Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. J Neurosci 16:4207–4221.

Boynton GM, Demb JB, Glover GH, Heeger DJ (1999) Neuronal basis of contrast discrimination. Vision Res 39:257–269.

Brainard DH (1997) The psychophysics toolbox. Spat Vis 10:433–436.

Buckner RL (1998) Event-related fMRI and the hemodynamic response. Hum Brain Mapp 6:373–377.

Campbell FW, Robson JG (1968) Application of Fourier analysis to the visibility of gratings. J Physiol (Lond) 197:551–566.

Carandini M, Movshon JA, Ferster D (1998) Pattern adaptation and cross-orientation interactions in the primary visual cortex. Neuropharmacology 37:501–511.

Dale AM, Buckner RL (1997) Selective averaging of rapidly presented individual trials using fMRI. Hum Brain Mapp 5:329–340.

Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN (1994) fMRI of human visual cortex. Nature [Erratum (1194) 370:106] 369:525.

Epstein R, Graham KS, Downing PE (2003) Viewpoint-specific scene representations in human parahippocampal cortex. Neuron 37:865–876.

Glover GH (1999) Deconvolution of impulse response in event-related BOLD fMRI. NeuroImage 9:416–429.

Greenlee MW, Georgeson MA, Magnussen S, Harris JP (1991) The time course of adaptation to spatial contrast. Vision Res 31:223–236.

Grill-Spector K, Malach R (2001) fMR-adaptation: a tool for studying the functional properties of human cortical neurons. Acta Psychol (Amst) 107:293–321.

Grill-Spector K, Kushnir T, Edelman S, Avidan G, Itzchak Y, Malach R (1999) Differential processing of objects under various viewing conditions in the human lateral occipital complex. Neuron 24:187–203.

Heeger DJ, Ress D (2002) What does fMRI tell us about neuronal activity? Nat Rev Neurosci 3:142–151.

Huettel SA, McCarthy G (2000) Evidence for a refractory period in the hemodynamic response to visual stimuli as measured by MRI. NeuroImage 11:547–553.

Kourtzi Z, Kanwisher N (2001) Representation of perceived object shape by the human lateral occipital complex. Science 293:1506–1509.

Levitt JB, Kiper DC, Movshon JA (1994) Receptive fields and functional architecture of macaque V2. J Neurophysiol 71:2517–2542.

Logothetis NK (2002) The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. Philos Trans R Soc Lond B Biol Sci 357:1003–1037.

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

Maddess T, McCourt ME, Blakeslee B, Cunningham RB (1988) Factors governing the adaptation of cells in area-17 of the cat visual cortex. Biol Cybern 59:229–236.

Miller KL, Luh WM, Liu TT, Martinez A, Obata T, Wong EC, Frank LR, Buxton RB (2001) Nonlinear temporal dynamics of the cerebral blood flow response. Hum Brain Mapp 13:1–12.

Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat Vis 10:437–442.

Ringach DL, Hawken MJ, Shapely R (1997) Dynamics of orientation tuning in macaque primary visual cortex. Nature 387:281–284.

Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and laminar dependence. J Neurosci 22:5639–5651.

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

Singh KD, Smith AT, Greenlee MW (2000) Spatiotemporal frequency and direction sensitivities of human visual areas measured using fMRI. NeuroImage 12:550–564.

Tootell RB, Hadjikhani NK, Vanduffel W, Liu AK, Mendola JD, Sereno MI, Dale AM (1998) Functional analysis of primary visual cortex (V1) in humans. Proc Natl Acad Sci USA 95:811–817.

Zheng D, LaMantia AS, Purves D (1991) Specialized vascularization of the primate visual cortex. J Neurosci 11:2622–2629.