

Neuronal correlates of perception in early visual cortex

David Ress¹ and David J. Heeger²

¹ Department of Psychology, Stanford University, 450 Serra Mall, Bldg. 420/400, Stanford, California 94309, USA

² Department of Psychology and Center for Neural Science, New York University, 6 Washington Place, 8th floor, New York, New York 10003, USA

Correspondence should be addressed to D.R. (ress@stanford.edu)

Published online 10 March 2003; doi:10.1038/nn1024

We used functional magnetic resonance imaging (fMRI) to measure activity in human early visual cortex (areas V1, V2 and V3) during a challenging contrast-detection task. Subjects attempted to detect the presence of slight contrast increments added to two kinds of background patterns. Behavioral responses were recorded so that the corresponding cortical activity could be grouped into the usual signal detection categories: hits, false alarms, misses and correct rejects. For both kinds of background patterns, the measured cortical activity was retinotopically specific. Hits and false alarms were associated with significantly more cortical activity than were correct rejects and misses. That false alarms evoked more activity than misses indicates that activity in early visual cortex corresponded to the subjects' percepts, rather than to the physically presented stimulus.

For more than 30 years, psychophysical studies of visual pattern discrimination have paralleled research on the neurophysiological response properties of neurons in the visual cortex^{1,2}. The prevailing view has been that psychophysical judgments about visual patterns are limited by neuronal signals in early visual cortical areas (such as primary visual cortex, V1). Signal detection theory has provided a theoretical framework for linking psychophysics and physiology. One testable property of this class of models of human pattern vision is that activity in early visual cortex should correspond to the subjects' percepts, even when those percepts are inaccurate.

The relationship between psychophysics and neurophysiology, as predicted by signal detection theory, can be studied using a contrast detection task (Fig. 1). On each trial, subjects were presented with one of two stimuli, either a background pattern presented alone or the same background with a low-contrast target pattern superimposed on it. Subjects pressed a button to indicate whether they believed the target was present or absent. Logically, there are four possible outcomes on a given trial: hits, when the observer correctly responds 'yes' on a target-present trial; correct rejects, when the observer correctly responds 'no' on a target-absent trial; false alarms, when the observer erroneously responds 'yes' on a target-absent trial; and misses, when the observer erroneously responds 'no' on a target-present trial. Because nearly all neurons in early visual cortex increase their activity monotonically with contrast²⁻⁴, target-present stimuli will, on average, evoke slightly greater neuronal activity than will target-absent stimuli. Neuronal responses vary, however, from one trial to the next, even when physically identical stimuli are presented repeatedly⁴⁻⁷. This variability in neuronal responses implies that a target-present stimulus can sometimes evoke less activity than a target-absent stimulus (Fig. 1, overlap between the two probability distributions). According to a simple model of the decision process, observers respond 'yes' when the neuronal activity exceeds a

fixed criterion (the vertical line in Fig. 1), and otherwise they respond 'no'. This criterion divides the two response distributions into four parts corresponding to the four possible outcomes. According to this model, we would expect the cortical activity averaged over many neurons and many trials of each outcome category to rank as follows: hits > false alarms > misses > correct rejects. This prediction is intuitive for the trials when the subject responds correctly (hits > correct rejects); cortical activity should be greater when the target contrast pattern is physically present in the stimulus. The prediction is counter-intuitive for the error trials (false alarms > misses); cortical activity now follows the subject's percept, which is the opposite of what is physically presented in the stimulus.

We used fMRI to measure activity in early visual cortex while subjects performed the threshold contrast-detection task. We observed that the average cortical activity within retinotopically predefined patches of cortex ranked as follows: responses to hits = false alarms > correct rejects = misses. Cortical activity during the error trials did indeed follow the subject's percept.

RESULTS

Subjects viewed a uniform gray field and continuously fixated on a small, high-contrast mark at its center while lying in the bore of the magnetic-resonance imaging scanner. Once every 2 s, a visual stimulus was displayed in an annulus around the fixation mark (Fig. 2) for 1 s, and a response period followed. On most of the trials, only a background pattern was presented; on the remaining (~1/6, randomly interleaved) trials, a low-contrast target grating was added to the background. Subjects pressed one of two buttons to indicate whether they thought the target was present or absent. Two different kinds of background patterns were used in separate experiments: plaid and noise. The plaid background provided a fixed stimulus configuration for which the only trial-to-trial variable was the presence or absence of the target. The noise background provided a

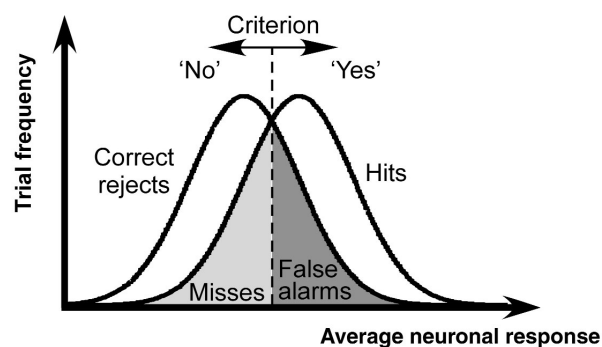


Fig. 1. An ideal-observer model of contrast detection. The observer makes his/her decision on each trial by comparing a noisy internal response (e.g., average firing rate of an appropriate subpopulation of neurons) with a fixed criterion. According to this model, the mean responses (across trials) for the 4 trial categories should rank: hits > false alarms > misses > correct rejects.

variable stimulus configuration, reducing the *a priori* information available to the subject and thereby making the task more difficult. Consequently, to maintain a fixed level of performance accuracy, target-contrast increments were considerably higher for the noise background (3.5–4.2%) than they were for the plaid background (0.6–0.9%).

fMRI data were collected in visual cortex during several thousand trials for each of four subjects and for both background patterns. The large number of trials was required to reliably measure the small fMRI signal changes (~0.1–0.2%) associated with the threshold-level stimulus contrast increments. Data were analyzed separately in visual areas V1, V2 and V3. (Data were also collected in areas V3A and hV4 in some, but not all, subjects and are shown in Table 1.) The analysis was restricted to the subregion of each visual area that corresponded retinotopically to the visual field location of the stimulus annulus. The trials were sorted into the four signal-detection categories: hits, false alarms, misses and correct rejects. Because the target pattern was presented infrequently, most of the trials (~70%) corresponded to correct rejects. Hence, fMRI activity levels associated with correct reject trials were taken as a baseline, and we calculated the differential activity associated with hits, misses and false alarms.

The main result was that cortical activity corresponded to trial category as follows: hits \approx false alarms > correct rejects \approx misses. This was evident in the fMRI time series acquired from individual visual cortical areas in individual subjects (Fig. 3a and b). It was also evident in a univariate measure of the fMRI response amplitude, which was computed by averaging the activity over a time window in the vicinity of the peak activity (Fig. 3c and d). Similar results were obtained using the peak fMRI amplitude as a measure, rather than the time-average in the vicinity of the peak (data not shown). When the data were combined across all subjects, the same ranking (hits \approx false alarms > correct rejects \approx misses) was again clearly evident (Fig. 4a and b).

To verify the statistical significance of these results, we compared the response amplitudes for the various trial categories, specifically to test the ranking predicted by signal-detection theory: hits > false alarms > misses > correct rejects. For completeness, we tested all possible two-way comparisons between the four categories, a total of six comparisons (hits > false alarms, hits > misses, and so on). We performed *t*-tests on response

amplitudes grouped only by trial category, that is, collected across subjects and visual areas. Four of the comparisons (hits > misses, hits > correct rejects, false alarms > misses, false alarms > correct rejects) were highly significant, as predicted by signal detection theory, for both the plaid-background data ($P \ll 0.001$) and for the noise-background data ($P \ll 0.001$). These effects were also significant in many individual visual areas in individual subjects (Table 1). Hits > false alarms was statistically significant only in the plaid-background data ($P = 0.015$), and not significant in the noise-background data ($P = 0.10$). Hence, there may have been small differences between the response amplitudes to hits and false alarms, as predicted by signal detection theory, but we could not resolve them with confidence. Misses > correct rejects was not significant in either data set ($P \approx 0.9$); in fact, this comparison showed the opposite trend (correct rejects > misses) as discussed further below.

We also performed a three-way ANOVA (by category, visual area and subject), allowing two-factor interactions. For both kinds of backgrounds, the ANOVA confirmed a main effect of trial category ($P \ll 0.001$). The only other significant main effect was by subject in the noise-background data, and there were significant interaction terms of subject \times category in both data sets; these effects may reflect individual differences in hemodynamic impulse response functions⁸.

The data from the plaid and noise backgrounds were similar, except that the peak response amplitudes for the noise background were larger (~40% on average) than those for the plaid background (Fig. 4, compare a and b), and the standard deviations were smaller. For the noise background, consequently, we were able to discern the differences between trial categories with considerably fewer trials per subject.

These results were retinotopically selective. We repeated our data analysis using the subregion of each visual area corresponding to a peripheral annulus (15–30° radius) well beyond the actual stimulus annulus. For the plaid background, response amplitudes were much smaller in the cortical representation of the periphery than the stimulus annulus for all




Fig. 2. Experimental stimuli: (a) plaid background; (b) plaid + vertical-grating target; (c) noise background; (d) noise + vertical-grating target. (Noise-background targets had randomized orientation and spatial phase.)

Table 1. Comparisons of fMRI response amplitudes.

Experiment	Subject	False alarms > misses					Hits > correct rejections				
		V1	V2	V3	V3A	V4	V1	V2	V3	V3A	V4
a) Plaid bkgd	DJH	0.033	0.030	0.108			0.263	0.006	0.000		
	BZL	0.119	0.015	0.081	0.119	0.195	0.000	0.022	0.001	0.000	0.001
	ACH	0.044	0.558	0.806			0.002	0.009	0.012		
	DBR	0.001	0.004	0.026	0.001	0.313	0.000	0.002	0.001	0.000	0.008
	All subjects	0.000	0.000	0.020	0.000	0.158	0.000	0.000	0.000	0.000	0.000
b) Noise bkgd	DJH	0.009	0.014	0.000			0.001	0.012	0.000		
	BZL	0.002	0.028	0.099	0.073	0.125	0.086	0.131	0.154	0.164	0.015
	AJN	0.030	0.089	0.005	0.047	0.048	0.001	0.013	0.000	0.001	0.000
	DBR	0.016	0.037	0.075	0.005	0.053	0.158	0.151	0.033	0.004	0.016
	All subjects	0.000	0.001	0.000	0.002	0.012	0.000	0.000	0.000	0.000	0.000

Numbers in table are null-hypothesis probabilities (*P*-values, one-tailed *t*-test) for false alarms > misses and hits > correct rejections, for each subject and visual area in the two experiments: plaid background and noise background. Significant values (*P* < 0.05) are in bold.

three visual areas, and for both hits and false alarms (*P* < 0.003; Fig. 4, compare a with c, and b with d). Similar results were obtained for the noise background (*P* < 0.02), except for false alarms in area V2 (*P* = 0.16).

The data were further analyzed to assess any correlation between the measured cortical activity and behavioral performance accuracy. A comparison of fMRI response amplitudes for correct (hits, correct rejections) versus incorrect (false alarms, misses) behavioral judgments showed no significant difference

(plaid background, *P* = 0.47; noise background, *P* = 0.078). There was a trend for response amplitudes to be higher for correct rejections than for misses, but it was weak (plaid background, *P* = 0.37; noise background, *P* = 0.034)⁹.

Because the noise background pattern was variable from trial to trial, we were concerned that the variability in the background pattern itself might bias the subjects' perceptual judgments, thereby confounding the interpretation of the results. A random-noise background can occasionally, by chance,

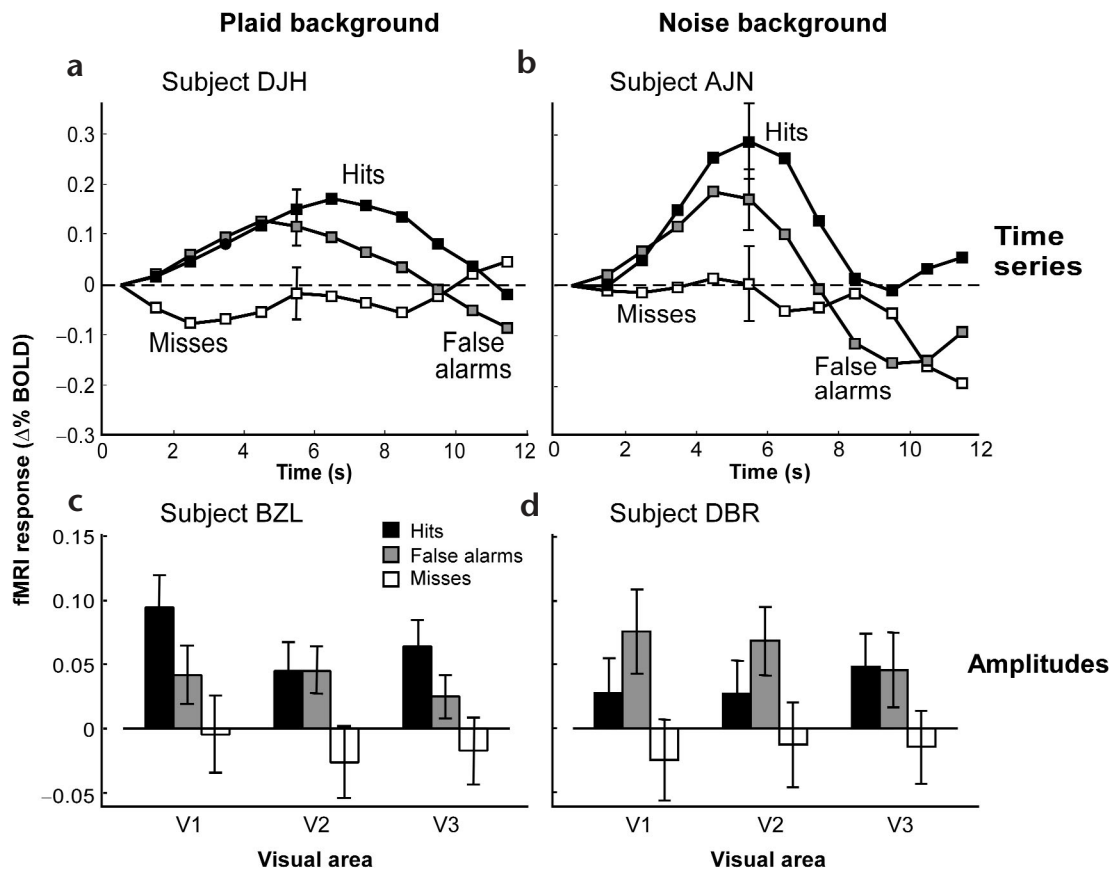


Fig. 3. Typical fMRI responses for individual subjects: (a, b) time series from V1; (c, d) response amplitudes; (a, c) plaid background; (b, d) noise background. Error bars are standard error of the mean (s.e.m.).

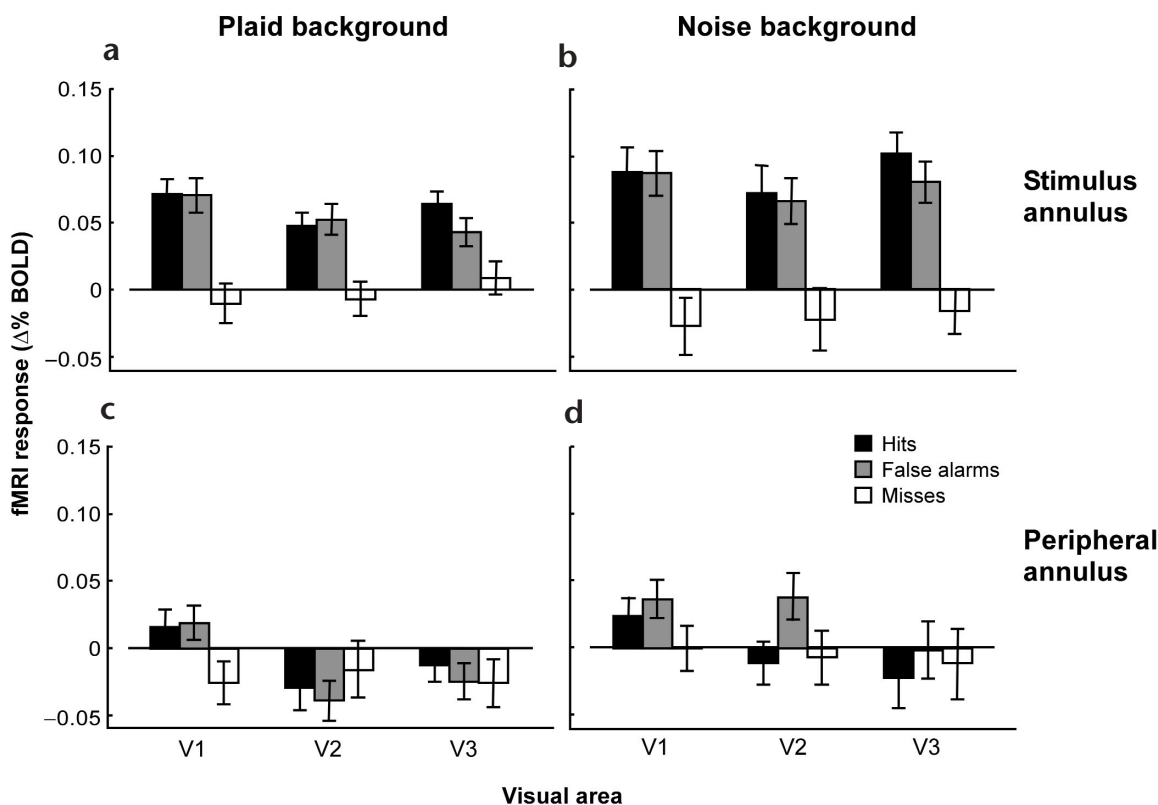


Fig. 4. fMRI response amplitudes averaged across subjects. (**a, b**) Response amplitudes from gray matter regions corresponding to the cortical representations of the stimulus annulus. (**c, d**) Response amplitudes from cortical representations of a peripheral region of the visual field. (**a, c**) Plaid background. (**b, d**) Noise background. Error bars are s.e.m.

resemble the target grating. Previous psychophysical studies have shown that under some circumstances, subjects show a tendency toward false alarms on trials in which the noise resembles the target^{10,11}. We made several choices in the design of our experiment (long stimulus presentation, multiple noise frames, randomized target orientation and phase) to avoid biasing subjects in this way. We tested for bias by analyzing the actual stimulus images that were presented on each trial, using methods similar to those used previously¹¹. There was no indication that the trial-to-trial variability in the noise background biased the subjects' perceptual judgments.

DISCUSSION

The central conclusion from these results is that activity in early visual cortex corresponded to the subjects' percepts, even when the percepts were the opposite of what was physically presented in the stimulus. We infer that trial-to-trial variability in the neuronal activity caused the trial-to-trial variability in perception. The trial-to-trial variability could result primarily from noise sources in bottom-up sensory processing, from noise sources in top-down (that is, non-sensory) processing, or from a combination of the two.

The measured signals levels are largely, but not entirely, consistent with the bottom-up sensory processing hypothesis offered by signal detection theory. The retinotopic specificity of our data supports this hypothesis, but there was little difference between the response amplitudes to hits and false alarms, and likewise between misses and correct rejections, as would be expected from the purely bottom-up interpretation.

Alternatively, performance might be limited by trial-to-trial variability in top-down signals. For example, subjects might perform the detection task by comparing the bottom-up visual signals with top-down signals that correspond to a working-memory representation of the target. The ability to balance the top-down and bottom-up inputs to visual cortex could contribute to or even dominate the variability in the reported percepts so that false alarms occur on those trials for which there are particularly strong working-memory representations, whereas misses correspond to weak memory representations. Further experiments will have to be performed specifically to test this (or a related) top-down interpretation of our results.

It is unlikely that our results were caused by a non-sensory neuronal process (such as arousal) associated with 'yes' responses. First, the results were retinotopically selective within early visual areas. Second, the results were stimulus dependent. Psychophysical contrast thresholds were higher for the noise background than the plaid background because the noise was higher contrast than the plaid and hence a more effective masker, and because there was more uncertainty with the noise background (the target orientation and spatial phase were randomized). fMRI response amplitudes were correspondingly larger in the noise-background experiment than in the plaid experiment.

Visual attention affects both behavioral performance and cortical activity^{9,12-16}, especially during threshold contrast detection⁹. Retinotopically localized activity is increased in early visual areas by attention¹³⁻¹⁵ even in the absence of a visual stimulus^{9,16,17}. In fact, previous work in our laboratory

showed that attention-related signals dominate fMRI measurements of cortical activity in early visual areas during a contrast detection task, and those attentional signals are correlated with performance accuracy so that cortical activity is greater, on average, for correct than for incorrect trials⁹.

The protocol in the current experiment had two features designed to control the effects of attention that had dominated our previous experiments. First, we used a rapid trial sequence that engaged subjects nearly continuously in the task. In our previous work, subjects were permitted to completely disengage their attention during the long inter-trial intervals. In our present experiments, the rapid pacing encouraged subjects to maintain their attention as steadily as possible. Second, we used an easily visible background pattern to minimize spatial uncertainty. In our previous work, subjects detected a low-contrast target on an otherwise blank (uniform gray) background so that there was uncertainty about the precise spatial location of the target, even after extensive practice. In the present experiments, the background pattern allowed subjects to allocate spatial attention to the correct location of the annulus on every trial.

In contrast to our previously published results, the present data show only a weak correlation between behavioral performance and cortical activity. A significant trend for such a correlation was evident only in the response amplitudes to correct rejects > misses. We believe that the experimental protocol of the present experiment reduced, but did not eliminate, the trial-to-trial variability in attention that dominated our previously published results. Further experiments will be needed to determine which of the differences between the two experimental protocols were critical for controlling the trial-to-trial variability in attention.

Signal detection in humans has been previously studied with event-related potentials. An early study measured evoked potentials over the occipital lobe while subjects performed contrast detection, and found greater evoked potential responses to hits than correct rejects, but no significant response to false alarms¹⁸. Particular attention has been paid to a transient evoked potential called the P3 or P300 (for example, see refs. 19–22). Under auditory stimulus conditions analogous to our experiments, the magnitude of the P3 exhibited dependence on trial category similar to what we have observed with fMRI^{19,20}. A common interpretation of the P3 is that it reflects a working-memory representation of the target that is used in the process of stimulus discrimination²¹. Another common interpretation of the P3 is that it reflects an 'oddball' effect, a response to the detection of infrequent targets²⁰. Further experiments will have to be performed to elucidate the relationship between the P3 and our fMRI measurements.

Previous fMRI experiments have found that a number of brain areas show greater responses during correct target detection (hits > misses, hits > correct rejects)^{23–27}. Some of these results might reflect neuronal processes similar to those that evoke the P3 waveform²⁶. However, none of these experiments were designed to measure the responses to false alarms or to assess the retinotopic specificity of the results, making a meaningful comparison to our present results difficult.

A recent fMRI study reports a result that is strikingly similar to our own, but measured in an area of visual cortex that is believed to subservise object recognition²⁸. In these experiments, fMRI responses were measured while subjects attempted to identify briefly presented complex objects. Trials were categorized as correct identifications, false identifications or missed identifications (analogous to our hits, false alarms and misses). Activity reflected the subjects' reported percepts, such that fMRI responses to correct identifications = false identifications > missed identifications.

Threshold-level signal detection has also been studied by measuring single- and multi-unit activity in animal models. One series of experiments shows that threshold level visual-motion signals are coded by the spike rates of neurons in cortical area MT, an area specialized for visual motion^{2,29–31}. Our results suggest an analogous process in human pattern vision, as early as V1. A second line of research has used backward masking to affect the detection of a visual target while recording single-unit activity in macaque frontal eye field neurons^{32,33}. The initial responses of these neurons (<100 ms after stimulus onset) corresponded to the actual presence or absence of the target, whereas later activity (100–300 ms after stimulus onset) corresponded to the monkey's behavioral judgments: responses were greater for false than for misses. Our results are analogous to this later phase (false alarms > misses), but in earlier visual areas. A third series of experiments recorded multi-unit activity in V1 while monkeys performed a figure-ground discrimination task³⁴. Once again, two phases of responses were evident, an early phase (<90 ms) that was directly related to the stimulus, and a later phase (100–240 ms) that was stronger when the monkey correctly performed the task (hits > misses). The later activity was attributed to feedback mechanisms. Indeed, several lines of investigation suggest that visual signals propagate rapidly from early visual cortex to other regions of the brain and back to visual cortex^{34–39}. The activity in early visual cortex associated with our detection task may likewise consist of an immediate response to the stimulus and a later feedback signal associated with the percept. If so, the sluggish hemodynamics of our fMRI measurements would have averaged this activity over time, yielding a superposition of both bottom-up sensory inputs and subsequent top-down feedback signals.

Our results suggest that early visual areas do more than encode raw sensory signals: they also participate in processing activities that correspond to a visual percept. The present results corroborate previous studies of ambiguous or bi-stable stimuli that show a correlation between the percept and activity in early visual cortex^{40–43}. The present results go further by showing that, during a threshold contrast detection task, both the percept and its corresponding cortical activity can be opposite to the physical stimulus. Thus, we conclude that perceptual errors are physically manifest by neuronal activity in early visual cortex.

METHODS

Magnetic resonance imaging. MR imaging was performed on a research-only GE 3T Signa scanner with a custom-designed dual surface coil. The experiments were undertaken with the written consent of each subject, and procedures were approved in advance by the Stanford Internal Review Board on Human Subjects Research. Each of five subjects participated in several MR scanning sessions: one to obtain a high-resolution, anatomical volume; one to functionally define the early, retinotopic visual areas including V1, V2, V3, V3A and hV4; and a variable number of sessions to measure fMRI responses in the various experimental conditions. Four of the subjects (DBR, DJH, BZL, ACH) participated in eight sessions devoted to the plaid-background task, and four of the subjects (DBR, DJH, BZL, AJN) participated in two sessions devoted to the noise-background task.

Each MR scanning session began by acquiring a set of anatomical images using a T1-weighted SPGR pulse sequence (TR = 10 ms, minimum TE, FA = 15°, 6 NEX, FOV = 220 mm, 4-mm slice thickness) in the same slices as the functional images. The eight slices were arranged obliquely, perpendicular to the calcarine sulcus, with the most caudal slice approximately tangent to the occipital pole. These inplane anatomical images were aligned to the high-resolution anatomical volume (acquired using a three-dimensional SPGR pulse sequence

and a head coil) of each subject's brain so that all MR images (across multiple scanning sessions) from a given subject were coregistered to an accuracy of ~ 1 mm⁴⁴.

Each fMRI scanning session included at least ten functional scans. Sessions began with a flickering-annulus reference scan described below. Next was a series of 9–11 contrast-detection scans with the target contrast set to that subject's detection threshold.

During each scan, a time series of fMRI volumes was acquired using a two-shot, T2*-sensitive, spiral-trajectory, gradient-recalled-echo pulse sequence^{45,46} (TE = 30 ms, TR = 500 ms, FA = 46°, FOV = 220 mm, effective inplane pixel size = 2.9 × 2.9 mm, 4-mm slice thickness).

A bite bar stabilized the subjects' heads. The fMRI images from each scan were visually inspected for head movements, which are evident as large steps or impulses in the time series. Fifteen out of ~480 scans showed evidence of head movements and were removed from further analysis.

Visual stimuli. Stimuli were presented on a flat-panel display (Multisync LCD 2000, NEC-Mitsubishi, Japan) placed within a Faraday box with a conducting glass front, positioned at the rear of the scanner bore. Subjects lay on their backs in the bore of the MR scanner and viewed the display through an angled mirror. Each subject had normal or corrected-to-normal vision.

During each contrast-detection fMRI scan, subjects performed 106 consecutive 2-s trials. Stimuli were presented in a near-foveal annulus (0.75–2.25° radius) and contrast modulated at 4 Hz. The plaid background was composed of a pair of diagonally oriented ($\pm 30^\circ$) sinusoidal gratings (10% contrast, 2.29 cpd), and the target pattern was a vertically oriented grating (0.6–0.8% contrast, 2 cpd). The noise background was composed of uniformly distributed intensity values (90% contrast). A sequence of four different noise patterns was presented during each 1 s stimulus, and the target was a superimposed grating with random orientation and spatial phase (3.5–4.2% contrast, 2 cpd). The target contrasts used during the fMRI experiments were individually chosen so that each subject would perform with an accuracy of $\sim 80\%$ correct ($d' \approx 1.5$). Before commencing fMRI scanning sessions, subjects practiced the task extensively until their performance stabilized.

Data analysis. Each fMRI time series was preprocessed by the following steps: (i) discarding the first 12 s of data to minimize transient magnetic-saturation effects, (ii) high-pass filtering the time series at each voxel to compensate for slow signal drift⁴⁷ and (iii) dividing each voxel's time series by its mean intensity. The resulting time series were averaged throughout the region of cortical gray matter corresponding to each pre-defined visual area (see below).

Time-series were calculated for each event category in the following manner. Individual trials corresponding to correct reject responses were treated as a baseline. The onset of the other three categories of trials initiated the corresponding event time series; this event time series was terminated by the onset of another non-correct-reject trial. Thus, an ensemble of variable-length event time series was extracted from the spatially averaged time series of each visual area. For each event category, we computed the mean time series and the standard error of the mean (s.e.m.) at each time point (Fig. 3a and b).

A univariate fMRI response amplitude was computed for each event by averaging the time series values from each visual area over a particular time window (typically 3–9 s after stimulus onset). The window was chosen to bracket the peak response determined from hemodynamic reference scans that were obtained separately for each subject.

The fMRI data were analyzed in each of five regions of interest (ROIs) corresponding to the V1, V2, V3, V3A and hV4 representations of the stimulus annulus in the cortical gray matter. These ROIs were defined, separately for each subject, in three steps. First, the visual areas were identified by measuring the polar angle component of the cortical retinotopic map^{48,49}. Second, we used an expanding-ring stimulus to identify the cortical representation of the stimulus annulus (0.5–2.25° radius) in each ROI. For our control measurements, we used similar methods to establish an ROI corresponding to the cortical representation of a peripheral annulus (15–30° radius). Third, the ROIs were further restricted, separately for each scanning session, according to a reference scan. During these reference scans, subjects held fixation while the display alter-

nated every 9 s between a uniform gray field and a contrast-reversing, high-contrast, plaid pattern within the stimulus annulus. The ROIs were restricted to regions that were strongly correlated with the stimulus alternations ($r > 0.5$; 0–6 s time lag).

Acknowledgments

The authors thank D. Nadell, P. Neri, A. Norcia, K. Shenoy, M. Silver and B. Wandell for comments. This research was supported by a National Eye Institute grant (R01-EY11794) and a grant from the Human Frontier Science Program (RG0070).

Competing interests statement

The authors declare that they have no competing financial interests.

RECEIVED 30 OCTOBER 2002; ACCEPTED 31 JANUARY 2003

- Graham, N. *Visual Pattern Analyzers* (ed. Broadbent, D.) (Oxford Univ. Press, New York, 1989).
- Parker, A.J. & Newsome, W.T. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* 21, 227–277 (1998).
- Carandini, M., Heeger, D.J. & Movshon, J.A. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* 17, 8621–8644 (1997).
- Geisler, W.S. & Albrecht, D.G. Visual cortex neurons in monkeys and cats: detection, discrimination and identification. *Vis. Neurosci.* 14, 897–919 (1997).
- Dean, A.F. The variability of discharge of simple cells in the cat striate cortex. *Exp. Brain Res.* 44, 437–440 (1981).
- Softky, W.R. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* 13, 334–350 (1993).
- Shadlen, M.N. & Newsome, W.T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18, 3870–3896 (1998).
- Aguirre, G.K., Zarahn, E. & D'Esposito, M. The variability of human, BOLD hemodynamic responses. *Neuroimage* 8, 360–369 (1998).
- Ress, D., Backus, B.T. & Heeger, D.J. Activity in primary visual cortex predicts performance in a visual detection task. *Nat. Neurosci.* 3, 940–945 (2000).
- Ahumada, A.J. Jr. Perceptual classification images from Vernier acuity masked by noise. *Perception* 26, 18 (1996).
- Solomon, J.A. Noise reveals visual mechanisms of detection and discrimination. *J. Vision* 2, 105–120 (2002).
- Pashler, H.E. *The Psychology of Attention* (MIT Press, Cambridge, Massachusetts, 1998).
- Tootell, R.B. *et al.* The retinotopy of visual spatial attention. *Neuron* 21, 1409–1422 (1998).
- Gandhi, S.P., Heeger, D.J. & Boynton, G.M. Spatial attention affects brain activity in human primary visual cortex. *Proc. Natl. Acad. Sci. USA* 96, 3314–3319 (1999).
- Brefczynski, J.A. & DeYoe, E.A. A physiological correlate of the 'spotlight' of visual attention. *Nat. Neurosci.* 2, 370–374 (1999).
- Kastner, S., Pinsk, M.A., De Weerd, P., Desimone, R. & Ungerleider, L.G. Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron* 22, 751–761 (1999).
- Chawla, D., Rees, G. & Friston, K.J. The physiological basis of attentional modulation in extrastriate visual areas. *Nat. Neurosci.* 2, 671–676 (1999).
- Campbell, F.W. & Kulikowski, J.J. An electrophysiological measure of the psychophysical contrast threshold. *J. Physiol.* 217, 54–55 (1971).
- Hillyard, S.A., Squires, K.C., Bauer, J.W. & Lindsay, P.H. Evoked potential correlates of auditory signal detection. *Science* 172, 1357–1360 (1971).
- Squires, K.C., Squires, N.K. & Hillyard, S.A. Vertex evoked potentials in a rating-scale detection task: relation to signal probability. *Behav. Biol.* 13, 21–34 (1975).
- Vogel, E.K., Luck, S.J. & Shapiro, K.L. Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *J. Exp. Psychol. Hum. Percept. Perform.* 24, 1656–1674 (1998).
- Hunter, M., Turner, A. & Fulham, W.R. Visual signal detection measured by event-related potentials. *Brain Cogn.* 46, 342–356 (2001).
- Shulman, G.L. *et al.* Areas involved in encoding and applying directional expectations to moving objects. *J. Neurosci.* 19, 9480–9496 (1999).
- Corbetta, M., Kincade, J.M., Ollinger, J.M., McAvoy, M.P. & Shulman, G.L. Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nat. Neurosci.* 3, 292–297 (2000).
- Shulman, G.L., Ollinger, J.M., Linenweber, M., Petersen, S.E. & Corbetta, M. Multiple neural correlates of detection in the human brain. *Proc. Natl. Acad. Sci. USA* 98, 313–318 (2001).
- Clark, V.P., Fannon, S., Lai, S., Benson, R. & Bauer, L. Responses to rare visual target and distractor stimuli using event-related fMRI. *J. Neurophysiol.* 83, 3133–3139 (2000).

27. Grill-Spector, K., Kushnir, T., Hendler, T. & Malach, R. The dynamics of object-selective activation correlate with recognition performance in humans. *Nat. Neurosci.* 3, 837–843 (2000).
28. Grill-Spector, K. *et al.* in *Attention and Performance XX* (eds. Kanwisher, N. & Duncan, J.) (Oxford Univ. Press, Oxford, in press).
29. Newsome, W.T., Britten, K.H. & Movshon, J.A. Neuronal correlates of a perceptual decision. *Nature* 341, 52–54 (1989).
30. Britten, K.H., Shadlen, M.N., Newsome, W.T. & Movshon, J.A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745–4765 (1992).
31. Shadlen, M.N., Britten, K.H., Newsome, W.T. & Movshon, J.A. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* 16, 1486–1510 (1996).
32. Thompson, K.G. The detection of visual signals by macaque frontal eye field during masking. *Nat. Neurosci.* 2, 44–49 (1999).
33. Thompson, K.G. Antecedents and correlates of visual detection and awareness in macaque prefrontal cortex. *Vis. Res.* 40, 1523–1538 (2000).
34. Super, H., Spekreijse, H. & Lamme, V.A. Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nat. Neurosci.* 4, 304–310 (2001).
35. Lamme, V.A., Super, H. & Spekreijse, H. Feedforward, horizontal and feedback processing in the visual cortex. *Curr. Opin. Neurobiol.* 8, 529–535 (1998).
36. Schmolesky, M.T. *et al.* Signal timing across the macaque visual system. *J. Neurophysiol.* 79, 3272–3278 (1998).
37. Liang, H., Ding, M., Nakamura, R. & Bressler, S.L. Causal influences in primate cerebral cortex during visual pattern discrimination. *Neuroreport* 11, 2875–2880 (2000).
38. Hupe, J.M. *et al.* Feedback connections act on the early part of the responses in monkey visual cortex. *J. Neurophysiol.* 85, 134–145 (2001).
39. Foxe, J.J. & Simpson, G.V. Flow of activation from V1 to frontal cortex in humans. A framework for defining 'early' visual processing. *Exp. Brain Res.* 142, 139–150 (2002).
40. Leopold, D.A. & Logothetis, N.K. Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* 379, 549–553 (1996).
41. Polonsky, A., Blake, R., Braun, J. & Heeger, D.J. Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nat. Neurosci.* 3, 1153–1159 (2000).
42. Blake, R. & Logothetis, N.K. Visual competition. *Nat. Rev. Neurosci.* 3, 13–21 (2002).
43. Tong, F. & Engel, S.A. Interocular rivalry revealed in the human cortical blind-spot representation. *Nature* 411, 195–199 (2001).
44. Nestares, O. & Heeger, D.J. Robust multi-resolution alignment of MRI brain volumes. *Magn. Reson. Med.* 43, 705–715 (2000).
45. Glover, G.H. & Lai, S. Self-navigated spiral fMRI: interleaved versus single-shot. *Magn. Reson. Med.* 39, 361–368 (1998).
46. Glover, G.H. Simple analytic spiral K-space algorithm. *Magn. Reson. Med.* 42, 412–415 (1999).
47. Smith, A.M. *et al.* Investigation of low frequency drift in fMRI signal. *Neuroimage* 9, 526–533 (1999).
48. Sereno, M.I. *et al.* Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268, 889–893 (1995).
49. Engel, S.A., Glover, G.H. & Wandell, B.A. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* 7, 181–192 (1997).