nature neuroscience

Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey

Alexander Maier[1], Melanie Wilke[1], Christopher Aura[1], Charles Zhu[2], Frank Q Ye[2] & David A Leopold[1,2]

The role of primary visual cortex (V1) in determining the contents of perception is controversial. Human functional magnetic resonance imaging (fMRI) studies of perceptual suppression have revealed a robust drop in V1 activity when a stimulus is subjectively invisible. In contrast, monkey single-unit recordings have failed to demonstrate such perception-locked changes in V1. To investigate the basis of this discrepancy, we measured both the blood oxygen level–dependent (BOLD) response and several electrophysiological signals in two behaving monkeys. We found that all signals were in good agreement during conventional stimulus presentation, showing strong visual modulation to presentation and removal of a stimulus. During perceptual suppression, however, only the BOLD response and the low-frequency local field potential (LFP) power showed decreases, whereas the spiking and high-frequency LFP power were unaffected. These results demonstrate that the coupling between the BOLD and electrophysiological signals in V1 is context dependent, with a marked dissociation occurring during perceptual suppression.

Perceptual suppression provides an intriguing puzzle for sensory neuroscientists: how is an unperceived stimulus represented in the brain? Binocular rivalry and related phenomena, in which salient visual patterns are rendered temporarily invisible to the observer, have been frequently employed to dissociate sensation from perception while monitoring brain activity[1]. These studies have provided insight into the nature of visual suppression, but have nonetheless failed to answer the fundamental question of what role V1 activity has in determining the visibility of a stimulus[2]. The main problem is that human and animal studies have reached nearly opposite conclusions. On the one hand, measurements of the BOLD signal in human fMRI have commonly found partial or even complete elimination of V1 responses to perceptually suppressed stimuli during rivalry[3–9], which suggests that V1 is important in determining whether a stimulus reaches subsequent processing stages. On the other hand, monkey single-unit studies found little or no change in activity in this area as a function of perceptual state, leading to the conclusion that perceptual suppression does not involve the perturbation of visual signals in the first cortical processing stage[10–13]. This discrepancy has been identified as a major impediment for understanding how V1 activity contributes to perception.

Until now, it has been impossible to isolate the potential reasons for the observed differences, as fMRI and electrophysiological studies have differed in their visual stimuli, behavioral procedures and species tested[3,14,15]. Here, we directly investigated the basis of these contradictory findings and further explored the role of V1 in visual perception by measuring both fMRI and electrophysiological responses of two trained monkeys experiencing perceptual suppression. Holding the stimulus procedure constant, we found that fMRI and neurophysiological signals in V1, although in good agreement during conventional stimulation, consistently diverged during perceptual suppression, with the most accurate reflection of stimulus visibility provided by BOLD responses.

RESULTS

The main objective of the study was to compare neural and fMRI responses under identical stimulus and perceptual conditions. To this end, we trained both monkeys to respond explicitly to the visibility of a salient target that was presented on a screen by depressing a lever whenever the target was present. We used generalized flash suppression (GFS)[16], where a salient target disappears on the abrupt onset of a surrounding field of dynamic dots, to reliably induce perceptual suppression (the stimulus sequence used in the present study is shown in **Fig. 1a**). Previous work in monkeys[13] and humans[16] demonstrated that such a stimulus reliably leads to all-or-none perceptual suppression in the absence of either local or interocular conflict. By adjusting stimulus parameters, such as the density of the dots, the minimum distance between the dots and the target, and the ocular configuration, we were able to affect the probability of the all-or-none target disappearance (see Methods). Suppression in the GFS procedure is related to that in motion-induced blindness[17], but is controlled using a temporal sequence resembling binocular rivalry flash suppression[18]. As with conventional binocular rivalry, GFS suppression is known to diminish the visual responses of extrastriate cortical neurons, but has little effect on V1 firing rates[10,11,19,20].

We used two general approaches to induce perceptual suppression. In the first, which we term perceptual report testing, we adjusted the

[1]Unit on Cognitive Neurophysiology and Imaging, Laboratory of Neuropsychology, National Institute of Mental Health and [2]Neurophysiology Imaging Facility, National Institute of Mental Health, National Institute of Neurological Disorders and Stroke, National Eye Institute, US National Institutes of Health, Department of Health and Human Services, 49 Convent Dr., B2J-45, MSC 4400, Bethesda, Maryland 20892, USA. Correspondence should be addressed to A.M. (maiera@mail.nih.gov).

Received 27 March; accepted 24 June; published online 24 August 2008; doi:10.1038/nn.2173

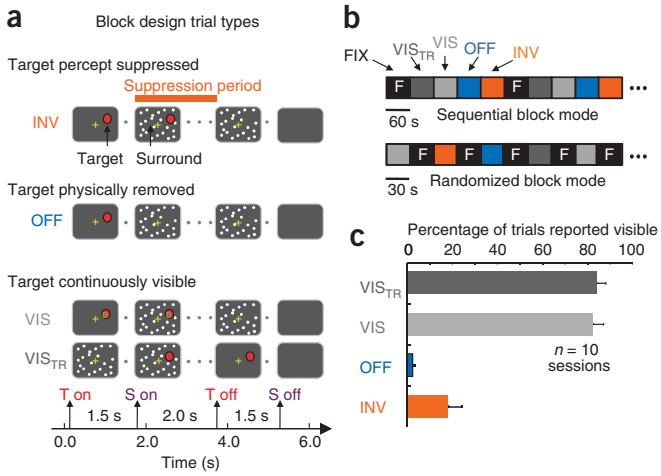

Figure 1 Generalized flash suppression protocol for the block design experiment. (**a**) Stimulus conditions. The parameters of the salient target and the random dot surround were adjusted to create five different conditions, with individual trials lasting 6 s (1-min blocks consisted of up to nine such trials). In the generalized flash suppression (INV) condition, the appearance of surrounding dots to both eyes consistently induced the monocular target to disappear (the time period during which the target stimulus disappears is indicated with a thick orange line). In the physical removal (OFF) condition, the target was physically extinguished on appearance of the dots, mimicking perceptual suppression. In the temporal reversal (VIS_{TR}) and binocular (VIS) conditions, the stimuli were adjusted by reversing the order of the target (red disk) and surround (moving random dots) or by presenting the target to both eyes, respectively, to ensure that the target did not disappear. Finally, in the fixation (FIX) condition (not depicted), the screen remained blank as the monkey fixated a small cross throughout the trial. (T on = target on, T off = target off, S on = surround on, S off = surround off). (**b**) The two modes of block design used for the fMRI experiments. In the sequential block design, the five stimulus conditions alternated in a fixed temporal order, keeping the number of stimulus conditions per run constant. In the randomized block design, each condition was presented in pseudorandomized order, each preceded by a blank screen fixation condition. (**c**) Psychophysical responses during the different conditions. The probabilities of perceptual suppression that are shown were based on ten psychophysical test sessions (five from each animal) and are plotted as the mean and s.e.m. over these sessions.

stimulus parameters at the beginning of each session according to the monkeys' psychophysical report such that the probability of disappearance for each trial presentation was roughly 0.5. This method, which was applied in a previous study[13], entails the sorting of neural responses by the monkey's report on a trial-by-trial basis. The second approach, which allowed direct comparison of fMRI and electrophysiological responses, was termed block design testing (**Fig. 1b**). In block design testing, which was used for the majority of the study, the visibility of the target was controlled over extended periods (blocks) by biasing the physical stimulus parameters (see Methods).

To isolate signal changes related to perceptual suppression, we compared responses during conditions in which the target was rendered either visible (VIS) or invisible (INV). Because of the existing discrepancy in the literature, we were particularly interested in whether activity during suppression (INV) would fall to the level of a control condition in which the stimulus condition was turned off when the surrounding dots appeared (OFF). This control condition is perceptually indistinguishable from suppression[16]. For both fMRI and electrophysiological testing, brain activation was measured relative to a stimulus-free fixation condition (FIX). Note that the monkeys were not required to confirm the disappearance during the block design testing (maintaining accurate perceptual report for truly ambiguous stimuli requires a large proportion of interspersed catch trials). Nonetheless, the effectiveness of the biased stimulus configurations was evaluated repeatedly in perceptual report testing. The behavioral results in the two animals verify that the manipulations of stimulus visibility were highly effective (see **Fig. 1c**) and matched the expectations of observed psychophysical results in a different group of monkeys[13].

In the block design, stimuli were grouped into 30- or 60-s epochs of repeated 6-s trials, during which the monkey was required to fixate a small cross. Blocks of the different conditions (INV, VIS, OFF, etc.) were presented either sequentially or in randomly interleaved order (**Fig. 1b**), with a total of four blocks of each type per scanning session. After successfully completed trials, the animal received a drop of apple juice and a short break (0.8–1.0s) in which it was free to briefly move its eyes about. Two animals participated in a total of 29 fMRI sessions and 41 electrophysiology sessions, which were conducted over the same time period, but on different days.

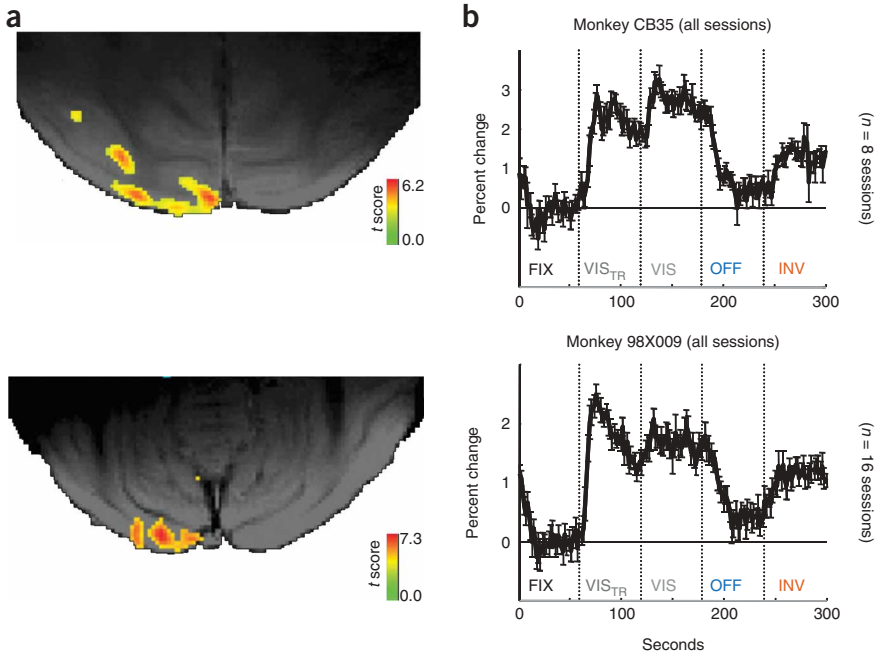

Figure 2 Modulation of BOLD responses during perceptual suppression in two monkeys. (**a**) Single-session examples of target-specific activation on axial slices (anterior is up, posterior is down) are shown. Colors represent the thresholded *t* score map corresponding to the statistical comparison between four repetitions of 30-s target presentation and four interleaved 30-s blocks of a blank screen (see Methods for the parameters used for the anatomical and functional magnetic resonance scans). (**b**) Mean BOLD responses over all sessions for both monkeys. Only voxels in V1 showing significant decreases in activity during the OFF period were included (see Methods for details). Mean ± s.e.m. over 8 and 16 sessions for Monkey CB35 and 98X009, respectively.

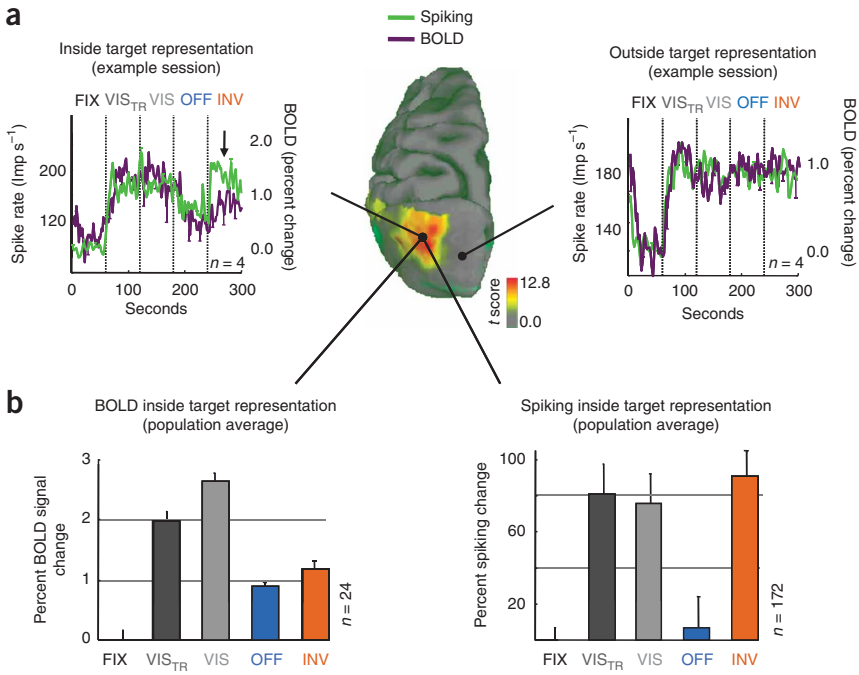

Figure 3 Divergence of V1 single-unit activity and fMRI BOLD response during perceptual suppression. (**a**) Single session examples of V1 BOLD responses and single neuron firing rates (in impulses per s) inside and outside of the target representation during the fMRI block design experiment. The colored region on the dorsal view of one hemisphere (anterior is up and medial is right) corresponds to the region of V1 that was activated by the target stimulus (data shown as *t* scores for a representative localizer experiment; Methods and **Supplementary Fig. 1**). In each of the panels, activity levels for the five different stimulus conditions are shown in the sequential block procedure (mean of four repetitions). Each trace represents the continuous activity level throughout 5 min of alternating 60-s stimulation blocks consisting of up to nine individual trials (vertical lines indicate the beginning and end of each stimulus block). Note that the BOLD and spiking activity drops in the OFF condition when the target is physically removed inside, but not outside, the target representation. The spiking and fMRI signals are in close correspondence, except for the GFS (INV) condition inside the target region. During this period, the BOLD signal showed perceptual modulation, whereas the spiking activity reflected the unchanged physical stimulation. Data from monkey CB35. Each plot is mean ± s.e.m. from four cycles of testing in one session. (**b**) Population average across both monkeys and all experiments. Left, mean BOLD response for all 24 scan sessions with time collapsed from the beginning to end of each 60-s block condition. Each bar represents the activation level in the region of interest as a function of condition. Note the drop in BOLD during the INV condition. Right, mean firing rate of all 172 recorded target-selective neurons, expressed as percent change to baseline activity (as assessed during the FIX condition). The VIS$_{TR}$ and VIS conditions resulted in statistically indistinguishable firing rates from the INV condition (multiple *t* tests, error bars represent s.e.m. between imaging sessions (left) and neurons (right)).

BOLD and spiking during perceptual suppression

We first localized the retinotopic region of V1 corresponding to the position of our recording chamber by functionally mapping the BOLD response to solid red disks in different positions (see Methods and **Supplementary Fig. 1** online). After finding the correspondence to the recording sites, we adopted this stimulus as a target for experimental testing. Our experiments revealed that perceptual suppression of the target (INV condition) strongly and consistently decreased BOLD responses compared with the continually visible (VIS) condition. Population data from two monkeys revealed that activity in the target responsive cortex dropped substantially during blocks of the invisible trials, compared with blocks of visible trials, where the target was continually perceived, even though the target was physically present in both cases (**Fig. 2**). Activity during the invisible trials was closest to the OFF control condition, where the target was physically removed from the screen.

The block design made it necessary to implement two distinctly different visible control conditions. Visibility was attained in one condition by binocular presentation of the target (VIS), and in the other by reversing the temporal order of the sequence (VIS$_{TR}$). The binocular presentation of the target in the VIS condition explains the slightly higher response than the one observed in the monocular VIS$_{TR}$ condition, consistent with previous work on binocular integration[21]. Nonetheless, this difference is small compared with the effects of perceptual suppression, where the target-responsive region showed activity that more closely resembled the effects of physically removing the stimulus. Note that in all cases, the target and surround stimuli were on for the same period of time.

Perceptual suppression was consistently observed in single sessions and in single voxels, and it was present when the order of the stimulus blocks was sequential or randomly interleaved (**Supplementary Fig. 2** online). These findings are consistent with a large number of human fMRI studies showing that BOLD signals in V1 are a good measure of perceptual visibility[3–9].

We next compared BOLD and single unit modulation during peceptual suppression (**Fig. 3**). We studied 172 well-isolated neurons whose responses were significantly modulated by the presentation of the target (*t* tests, $P < 0.05$, out of 318 total neurons sampled, see Methods). The monkey subjects, the stimulus and the block-design protocol were identical to the fMRI testing. In sharp contrast to the modulation of the BOLD signal, neurons in the target-responsive portion of cortex showed no significant change in their mean firing during perceptual suppression ($P = 0.98$, *t* test; **Fig. 3**). Notably, the BOLD and spiking responses in V1, tested in the same patch of cortex in the same monkey subjects, specifically diverged during perceptual suppression, with the BOLD signals, but not the spiking responses, following the percept. This result serves to reconcile previous single-unit studies in monkeys[10–13] with fMRI studies in humans[3–9] and demonstrates that neither the species nor the stimulation protocol formed the basis of the discrepancy, which instead arose from the nature of the measured signals themselves.

LFP modulation, perceptual suppression and the block design

We next explored whether modulation of the LFP signal might reflect the BOLD signal more closely than the spiking and therein match the monkey's perceptual state[22]. Recent work using the perceptual report procedure in GFS revealed that the power of lower-frequency LFP components (<30 Hz) substantially decreased when the monkey reported the disappearance of an ambiguous stimulus[13]. We thus investigated whether LFP modulation was present in our data by first dividing the LFP into two frequency bands, low (5–30 Hz) and high

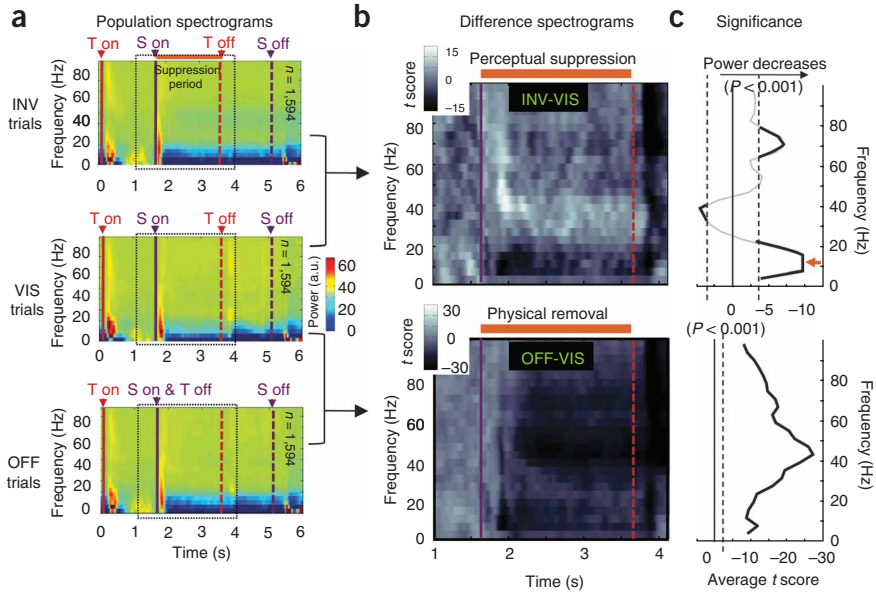

Figure 4 Spectral analysis of LFP signals obtained during suppression and control trials. (**a**) Time-frequency plots of the INV, VIS and OFF conditions, featuring activity changes in the suppression period (following surround onset). Each panel depicts the average spectrogram for an entire trial period, with magnitude changes relative to the 500-ms period preceding the surround onset, in which the target alone was present in each condition. Population data are shown over all channels and over all recording sessions (successfully completed trials only). All relevant stimulus events are marked with dashed lines (T = target and S = surround). Note that although all of the conditions showed a drop in low-frequency power following the surround onset, the drop for the INV condition during the suppression period was larger than the corresponding drop for the VIS condition and closely resembled physical removal (OFF). (**b**) Statistical time-frequency analysis of perceptual suppression versus physical removal (for the time period indicated by a dashed square in **a**). The upper plot (comparing the INV and VIS conditions) shows a decrease during perceptual suppression that was limited to the low frequencies, whereas the lower plot (comparing the OFF and VIS conditions) shows a large, broadband decrease when the target was physically removed. Note that these two conditions are nearly identical perceptually. *t* values are indicated by gray scale values shown in inset. (**c**) Average *t* score as a function of frequency for the entire period of perceptual suppression (and physical removal). The threshold for statistical significance ($P < 0.001$) is indicated by a dashed line.

(30–90 Hz), and then computing power changes in the different bands over blocks (**Supplementary Fig. 3** online). We reasoned that, given the robust BOLD changes, there might be low-frequency modulation during blocks in which the target was suppressed compared with those in which it was continuously visible. In contrast to our expectations, we found no significant decrease in the mean power for either the low- or high-frequency bands during blocks of perceptual suppression ($P = 0.69$ and $P = 0.68$, respectively, *t* test, 95 channels). In fact, we were unable to find any electrophysiological signal that, over the course of an entire block, correlated with perceptual suppression.

To explore the basis of this negative finding, we next sought to reproduce findings from our previous study, in which the LFP had reflected stimulus visibility in V1. That study used the perceptual report procedure (in different animals) and found that purely perceptual signals were reflected in the modulation of the low-frequency LFP[13]. We trained both monkeys to report the visibility of target stimuli on trials that were all identical, but in which the target had a 0.5 probability of being suppressed (see Methods).

In agreement with the previous results, we found clear and consistent power decreases in the low-frequency LFP during suppressed trials (**Supplementary Fig. 4** online). We then asked whether perceptual modulation had, in fact, been present in the block design trials, but had been diluted because of the relatively small proportion of time in the block in which the target was suppressed. To examine this possibility, we conducted a trial-by-trial analysis of the block design data, beginning with a time-frequency analysis.

We created spectrograms of the time course of power changes in different frequency bands within trials of different block types (**Fig. 4a**). We focused on the 2-s period following surround onset corresponding to the period of perceptual suppression in the INV trials, or lack thereof in the VIS trials. We evaluated which frequency components were affected by suppression (INV-VIS) and compared the effects of suppression to those of physical removal (OFF-VIS), plotting the *t* score of the

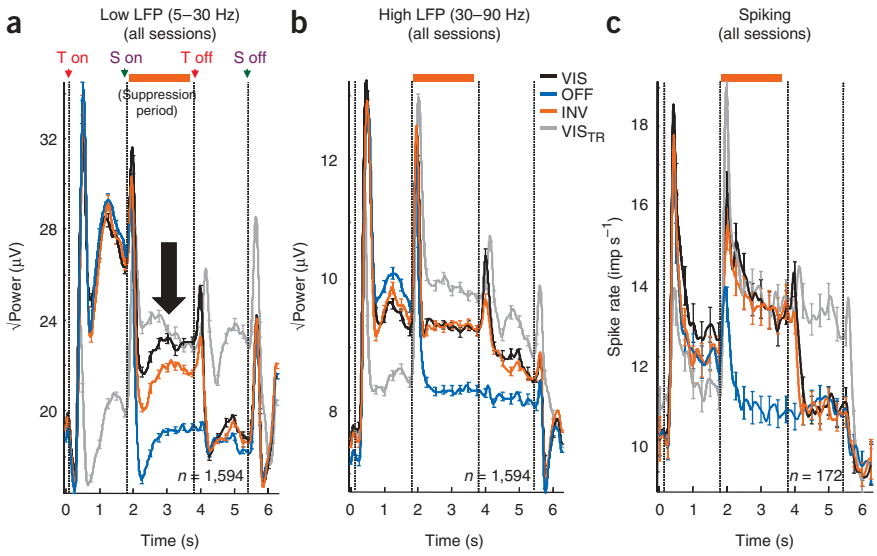

Figure 5 Population average of band-limited power (BLP) and spiking time courses for different experimental conditions. (**a**) Grand average of low-frequency (5–30 Hz) BLP over time as a function of experimental condition (all channels, all sessions, both monkeys). All relevant events are indicated with dashed lines (T = target and S = surround). The analysis time window (during which target was perceptually suppressed during the INV condition) is indicated by the orange bar. (**b**) Grand average of high-frequency (30–80 Hz) BLP, same conventions as in **a**. (**c**) Spiking density function of all units recorded during the neurophysiological experiments. In all cases, data are convolved with a 50-ms s.d. Gaussian kernel. Error bars are s.e.m.

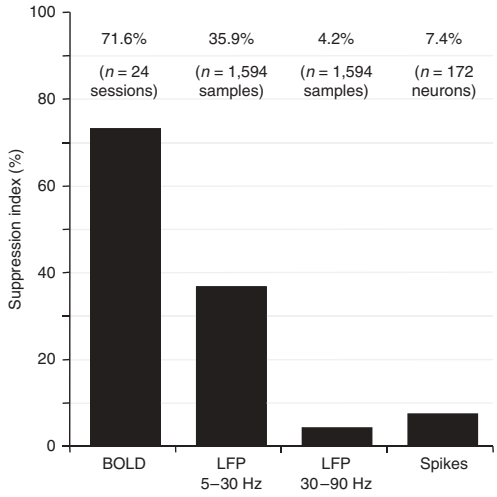

Figure 6 Summary of perceptual modulation in the BOLD response and in each of the electrophysiological signals (as computed from the raw data of both monkeys shown in **Figs. 3b** and **5**). The Suppression index corresponds to the percent of signal drop during the INV condition compared with the OFF condition, both relative to the VIS condition.

activity changes under the two conditions (**Fig. 4b,c**). Note that the physical removal of the visual target produced a large-amplitude broadband decrease in power, but that perceptual suppression did not (**Supplementary Fig. 5** online). Consistent with the perceptual report procedure, the strongest effect of suppression was a significant decline in the power of low frequencies ($P < 0.001$; **Fig. 4c**), matching the results obtained from the reporting monkey described above.

We examined the high (30–90 Hz) and low (5–30 Hz) frequency ranges, as well as spike rates, for each testing condition (the magnitude of low-frequency power modulation can be seen more clearly in **Fig. 5**). This analysis revealed that the perceptual suppression condition (orange) deviated significantly from ($P < 0.001$) the visible condition (black) in the low-frequency LFP (black arrow), but not in high-frequency LFP or in spiking ($P = 0.99$ and $P = 0.90$, respectively). In fact, the only detectable electrophysiological difference between the visible and invisible trials was this change in low-frequency LFP power. These data are summarized in **Figure 6**, which shows the suppression index, comparing the activity change during perceptual suppression to that measured during the physical removal of the target. Of these signals, it was the BOLD and low-frequency LFP that showed substantial declines relative to the control conditions, therein reflecting the state of perceptual suppression.

DISCUSSION

These findings help to resolve a long-standing discrepancy between human fMRI and monkey neurophysiology regarding the role of V1 in determining whether a stimulus is visible. The outcome of the combined fMRI/electrophysiological approach demonstrates that the different conclusions reached by human fMRI and monkey electrophysiology were not the results of either species or procedural differences, but can instead be attributed to the nature of the signals that were measured. Although consistent during conventional stimulation, the BOLD and electrophysiological responses diverged markedly during perceptual suppression.

It is interesting to consider that the BOLD fMRI signal, arguably the furthest removed from neural processing, provided the most reliable measure of the perceptual state, whereas the action potential firing rate provided the least. Why might this be the case? There are several possible explanations. First, one cannot entirely rule out an electrophysiological sampling bias that consistently missed a subpopulation of neurons, perhaps because of their size, that carries the perceptual signal in V1 (see ref. 23 for neuron-type specific modulation effects). Perceptual modulation in small, infrequently sampled neurons, such as interneurons bearing a close relationship to vascular control[24], might elicit a prominent BOLD response. Another possibility is that perceptual suppression results in a temporary and spatially localized mode of cortical processing in which both inhibition and excitation are decreased, but remain in the same balance. This scenario could produce a minimal or no change in the spiking of neurons, but would provide a temporary relief of the metabolic and vascular demands, leading to a decreased hemodynamic response. Finally, it is also possible that seemingly unimportant modulatory signals distributed over a large population of neurons would escape the notice of electrophysiological analysis, but would be effectively registered in the BOLD response. Further experiments are required to address these and other possibilities.

A potentially attractive explanation of our findings might be that input to V1 from structures other than the LGN is disrupted during perceptual suppression and that such disruption is better reflected in the BOLD changes than in neural firing. Considerable evidence suggests that synaptic activity (stemming from afferent input and local intracortical processing) might lead to BOLD responses, even in the absence of firing rate changes[25–27]. It has long been hypothesized that recurrent input from extrastriate cortical areas into V1 accompanies selective attention and stimulus awareness[28,29]. In this context, is also noteworthy that a similar mismatch between fMRI and single-unit physiology has been previously identified regarding the modulation of V1 by visual attention[30], suggesting that the dissociation observed here may arise in conditions other than perceptual suppression. Recent work demonstrating laminar differences in the suppressive modulation of synaptic currents in V1 (ref. 31) shows promise for providing a deeper understanding of both the neural mechanisms of perceptual suppression and the enigmatic relationship between the fMRI signal and neural activity.

The lack of LFP power modulation over blocks of perceptual suppression was unexpected, given that a clear LFP power modulation was seen in the trial-based analysis of the same data, along with the clear BOLD modulation in the block-based analysis. This result appears to indicate that the BOLD response is not a simple reflection of integrated neural activity over time. This negative finding may owe, in part, to a low signal-to-noise ratio resulting from a dilution over large time windows. However, it is interesting to consider that the robust suppression that we observed in the BOLD response might be shaped by only a subset of neural events, which are punctuated and context specific. That would have profound implications, as there could be no unique hemodynamic transfer function to serve as a time-invariant convolution filter to translate neural and BOLD signals, as is commonly assumed. For example, neural modulation associated with the stimulus-driven events may have a proportionally stronger contribution to the hemodynamic response than do neural fluctuations occurring between trials. At the same time, recent work relating spontaneous neural activity in V1 to BOLD fluctuations suggests that endogenous activity can also have a strong effect on the fMRI response[32].

Our results demonstrate that the very same signals that correlate strongly with the BOLD signal in one context (physical stimulus removal) fail to do so in another (perceptual suppression). The absence of a fixed relationship between the different neural signals is illustrated conceptually in **Figure 7**. When a target stimulus eliciting a tonic neural response is physically removed, all measured signals are consistent, showing corresponding decreases in amplitude. On the other hand,

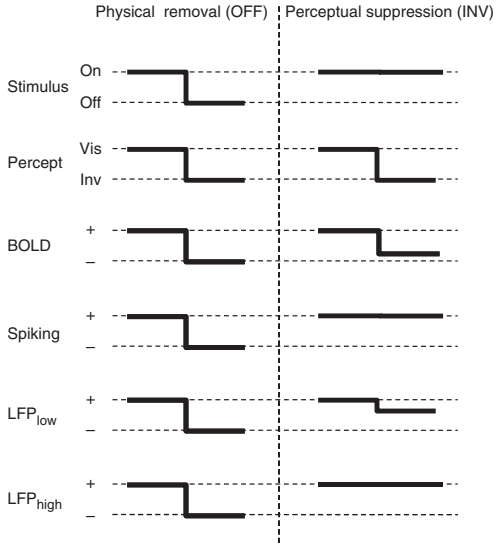

Figure 7 Schematic illustration of main results. From top to bottom, each line represents the state of presentation (stimulus either on or off), the reported percept of the subject (stimulus visible or invisible) and the various measures of neuronal activity in primary visual cortex (high or low activity). The left column represents the OFF condition. As shown in **Figures 2–5**, all measures of neuronal activity, including the fMRI BOLD response, showed a decline in signal when the stimulus was both physically removed and perceptually disappeared. The right column represents our finding for the INV condition, where the percept was dissociated from the physical stimulus in the form of perceptual suppression. In that case, the spiking activity (and high frequency LFP) maintained its activity level, reflecting the continually present stimulus, whereas the fMRI response and (to a lesser extent) the low-frequency LFP reflected the perceptual disappearance. Thus the relation between the fMRI and neural signals differs under the two conditions.

during perceptual suppression, the various responses become dissociated and the BOLD signal no longer matches the neurophysiological signals. Such context dependency might serve to explain discrepancies among previous studies investigating the neural basis of the BOLD signal[25–27,33,34]. Gaining a deeper understanding of the neural determinants of the BOLD fMRI signal remains an important challenge for systems neuroscience. The parallax provided by neural and hemodynamic signals, each providing a window onto different brain function, may eventually give mechanistic insights that neither technique could provide alone.

METHODS

Subjects and testing. Two healthy adult male monkeys (*Macaca mulatta*) were used in this study. All procedures followed US National Institutes of Health guidelines, were approved by the Animal Care and Use Committee of the US National Institutes of Health (National Institute of Mental Health) and were conducted with great care for the comfort and well being of the animals. Each monkey had an magnetic resonance–compatible recording well implanted over V1, which allowed for fMRI and neurophysiological testing on subsequent days. A total of 41 multi-electrode recording sessions (27 and 14 sessions in Monkeys CB35 and 98X009, respectively) and 29 fMRI sessions (8 and 16 sessions, respectively, as well as 5 additional sessions with randomly ordered blocks) were collected.

Stimulus and task design. All sessions were carried out while the animals were awake and performing a task, either fixation only or reporting target visibility. In the block design task, visual stimuli were presented in 30- or 60-s epochs of repeated 6-s trials of a given type. Animals were required to fixate a small spot until the completion of each trial to obtain an apple juice reward. Each stimulus type consisted of a target pattern (a bright red disk, with a diameter of 4° of visual angle (dva) for the majority of the experiments, was presented to the parafovea in one quadrant), and a surrounding field of randomly moving dots (200–500 total dots of 0.5° diameter, covering an area of 40° × 30°). The random dots never approached within half a degree of the target stimulus. In the suppression condition (INV), the stimulus parameters were adjusted such that the presentation of the surround pattern resulted in perceptual suppression of the target. In two 'visible' control conditions (VIS and VIS$_{TR}$), the parameters were adjusted such that the target almost never disappeared. In the 'invisible' control condition (OFF), the target was physically removed from the screen on appearance of the dots. Blocks were either presented in a repeated temporal sequence or randomly interleaved. In other testing, animals responded according to the visibility of the target by pulling a lever.

In most sessions, animals were tested in a block protocol consisting of 60-s blocks of repeated presentations of the five different stimulus types (FIX, VIS$_{TR}$, VIS, OFF and INV). These blocks were presented either sequentially or in pseudorandom order. In each block, the animals were required to fixate for the entire 6-s presentation of a particular stimulus condition, which was repeated until the end of the block. The quality of fixation was continually monitored. Eye movements exceeding 1–2° away from the fixation spot caused the trial to abort. Note that because the monkeys had been trained to fixate within 0.8° of the fixation spot, they kept their eyes in that region for the large majority of the time (**Supplementary Figs. 6–9** online). Following each successful trial, the animal received a juice reward, accompanied by a short break (0.8–1 s) in which the monkey was free to move its eyes about. There were typically 6–8 completed trials each minute, depending on how well the animal acquired and maintained fixation (no significant performance differences were found between the VIS, VIS$_{TR}$, OFF and INV conditions in either animal).

Differential stimulation to the two eyes was achieved inside and outside the scanner by fixing anaglyph filters in front of the monkeys' eyes and adjusting the stimulus color accordingly (see **Supplementary Methods** online for additional details).

Behavioral protocol. The experiments started after the monkeys were able to hold and maintain fixation on more than 95% of trials. Each session started with a brief calibration procedure, during which the monkeys were presented with a small (0.1–0.25 dva) fixation spot at one of nine positions on the screen. After the monkeys acquired and kept fixation for 1–3 s, juice reward was delivered and a new trial began. Following this procedure, we started the experimental task, in which the monkeys had to fixate a central spot on the screen in a 1–4-dva window for up to 10 s to receive reward. If a monkey broke fixation, the trial was aborted and re-initiated after a short delay of 100–800 ms.

We carried out psychophysical testing in several separate test sessions outside of the scanner. Stimulus disappearance was reported using custom-made levers mounted inside the primate chair. One monkey had previously been trained to report the physical and subjective removal of the target stimulus using a variety of catch trials to ensure truthful perceptual report. The other monkey was trained to do so at the end of the experimental sessions (for a more detailed description of the methods used for behavioral training and testing of GFS, see ref. 13). Stimulation parameters in the physiology testing booths matched those used in the scanner.

Surgery. Monkeys were implanted under isoflurane anesthesia (1.5–2%) with custom-designed and fabricated fiberglass headposts, which served to immobilize the head during testing. Each monkey received a scleral search coil, following a standard implantation procedure[35]. In a subsequent surgery, a recording chamber was implanted over V1 using frameless stereotaxy guided by high resolution anatomical magnetic resonance images (Brainsight, Rogue Research), and a craniotomy was performed[35]. Animals received antibiotics and analgesics post-operation.

fMRI scanning. Structural and functional images were acquired in a 4.7-T, 60-cm vertical scanner (Bruker Biospec) equipped with a Siemens AC44 gradient coil (40 mT m^{-1}, $<$200 μs). Magnetic resonance–compatible primate chairs were constructed and machined using plastic materials and with a minimum of brass and aluminum parts. Monkeys were prevented from performing excessive jaw movements by a chin rest that was mounted to the top of the chair. Juice was delivered using an air-pressurized juicer device[36]

through custom-made brass mouth pieces. Exposure to scanner noise was reduced by means of custom-formed ear plugs.

Monkeys were scanned over a period of 4 months (and two sessions of confirmatory retesting 12 months later) using two custom-built surface coils situated over cortical area V1. The animals had previously been acclimated to the scanner testing environment and worked continuously for up to 4 h. Electrophysiological testing was conducted during the same testing period on different days using the same animals, primate chairs, stimuli and behavioral testing parameters, in designated electrophysiology booths.

Each session began with a localizer scan, in which the retinotopic area corresponding to the salient target was identified. This localizer consisted of 60-s blocks of trials alternating between presentation of the target, the surrounding dots or a fixation cross presented alone. Analysis of the main experiment was restricted to those voxels that showed significant decreases when the target was physically removed in the context of the block design (in the OFF condition).

Transmit-and-receive radio frequency coils had dimensions ranging from 33 × 33 to 37 × 120 mm and were placed adjacent to the scalp of the animals before they entered the magnet bore. The size and position of the coils were optimized to achieve maximal sensitivity in the posterior portion of area V1 in the occipital lobe. At the beginning of each session, active shimming was achieved either by manual manipulation of the gain of the shim coils, by using the FASTMAP[37] procedure or by a combination of the two. To assess functional cortical activation, a single shot gradient echo-planar imaging[38] sequence was used with a repetition time between 2.0 and 2.5 s and an echo time between 30 and 35 ms. Up to ten axial slices were collected with a field of view ranging from 96 × 113 mm to 128 ×128 mm and a slice thickness of 1.25–2 mm. The in-plane resolution of the functional images was 1.5 × 1.5 mm for the majority of the scans.

fMRI analysis. Data analysis was carried out on a PC running the Windows operating system and custom-written software using MATLAB (MathWorks) as well as the AFNI/SUMA software package[39]. To analyze fMRI data, raw images were first converted from the generic BRUKER into the common AFNI data format. fMRI in monkeys is prone to image distortions arising from movements and postural adjustments. We minimized distortions over time by excluding time points for which ghosting or movement led to irreparable image distortion, and applying the AFNI function 3dWarpDrive to each time point, which applied an affine transformation to each time point on a slice by slice basis to match a reference slice. These measures substantially improved the quality of the spatial mapping, as well as the BOLD time courses.

To analyze functional activation, we converted the adjusted images of each scan into units of percent change by subtracting each run's mean value from the time series of each voxel, and subsequently dividing by the same value, on a voxel-by-voxel basis. Note that for presentation purposes, we later subtracted the mean activity during the fixation condition from activity in the other conditions, defining it to be our baseline for comparing experimental conditions). To correct for slow drifts in the magnetic resonance signal that were unrelated to the task, we high-pass filtered each scan using a bi-directional second-order Butterworth filter with a cutoff frequency of 0.0025 Hz.

A region of interest was defined as the target stimulus–responsive region of primary visual cortex that showed significant activity differences for target versus nontarget presentation conditions defined as any voxel in the anatomical limits of V1 exhibiting a t value more than two s.d. away from the average of the t score maps (**Supplementary Fig. 1**).

Neurophysiological recordings. Extracellular single-unit and LFP recordings were carried out using sharp insulated microelectrodes and multi-contact transcortical electrodes. Single-unit activity and LFPs were recorded from the primary visual cortex of both animals. Recordings were carried out inside a radio frequency–shielded booth that was also used for behavioral testing. In all cases voltages were measured against a local reference that was close to the electrode contacts (that is, a stainless steel guide tube or the hypodermic metallic shaft surrounding the multi-contact electrodes). Recording electrodes consisted of both single channel microelectrodes (Thomas Recording GmbH) as well as 16- and 24-multicontact contact electrodes with an intercontact spacing of 150 μm and 100μm, respectively (Neurotrack).

Single-unit and LFP activity was collected using both electrode types with the MAP recording system (Plexon). The V1 sites were located dorsally, several millimeters posterior to the lunate sulcus, and covered the parafoveal region close to the vertical meridian.

Single units were isolated and characterized in terms of their basic response characteristics using a custom-written program for receptive field estimation. Multiunit activity, in the form of voltage spikes exceeding a manually set threshold, were collected and digitized by the MAP recording system. Single-unit impulses were derived from the multiunit data by using a commercially available spike sorting program (Plexon). The quality of isolation was assessed and rated by two investigators, and only units with perfectly (that is, completely unambiguous) isolated clusters were included in the analysis. The LFPs (measured as voltage fluctuations between 1 and 100 Hz) were collected simultaneously using the same system and digitized at 1 kHz. Notably, during all of the recording sessions, the monkeys' task was similar in every respect to that inside the scanner. All behavioral and stimulus events were encoded and recorded together with the neuronal signal on a separate channel to align the data during *post hoc* analysis.

Neuronal data analysis. All neurophysiological data was processed and analyzed using custom-written code for MATLAB. Single-unit spiking data was converted into spike-density functions with a sample rate of 1 kHz by replacing each spike time with a Gaussian kernel. LFPs were resampled at 1 kHz and converted into microvolts as a function of time. Spectrograms were computed using the Fast Fourier Transform with a running window size of 256 ms and an overlap of 255 ms or the multitaper method (CHRONUX toolbox for Matlab, http://www.chronux.org/) using similar parameters. Both techniques yielded highly similar results.

In addition, we subdivided the LFP data into different bands in the frequency domain using a second-order, bi-directional, zero-phase Chebyshev type-1 filter. The resulting band-limited signals were full-wave rectified by taking their absolute value, and resampled at 200 Hz. Rectifying the band-limited voltage fluctuations results in a measure of time-varying amplitude, or signal power (in actuality, the square root of the power), in each frequency band. This band-limited power (BLP) is roughly equivalent to averaging several adjacent rows of a spectrogram (for an extended discussion, see ref. 40). For conditions in which percent change above baseline was computed, the baseline was taken to be the mean firing rate or mean band-limited LFP power measured during the fixation period.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We would like to thank G. Dold, D. Ide, N. Nichols and T. Talbot, as well as K. Smith, N. Phipps and J. Yu for technical assistance. We also thank H. Merkle for extensive guidance on the design and fabrication of radio frequency coils, R. Cox for assistance with magnetic resonance image alignment, D. Sheinberg for help with the stimulus software, K. King and C. Brewer for auditory testing and ear plug manufacture, W. Vinje for help with the multi-contact electrodes, K. Tanji, A.H. Bell and Z. Saad for help with the fMRI analysis, and S. Guderian, M. Schmid and K.-M. Mueller for discussions. This work was supported by the Intramural Research Programs of the National Institute of Mental Health, the National Institute of Neurological Disorders and Stroke and the National Eye Institute.

AUTHOR CONTRIBUTIONS

A.M., M.W. and D.A.L. designed the experiments. A.M., M.W., C.A., C.Z., F.Q.Y. and D.A.L. contributed to the fMRI experiments. A.M., M.W. and C.A. carried out the electrophysiological testing, and M.W. collected the psychophysical data. A.M. analyzed the fMRI and electrophysiological data. A.M. and D.A.L. wrote the paper.

Published online at http://www.nature.com/natureneuroscience/
Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions/

1. Blake, R. & Logothetis, N.K. Visual competition. *Nat. Rev. Neurosci.* **3**, 13–21 (2002).
2. Crick, F. & Koch, C. Are we aware of neural activity in primary visual cortex? *Nature* **375**, 121–123 (1995).

3. 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).
4. Tong, F. & Engel, S.A. Interocular rivalry revealed in the human cortical blind-spot representation. *Nature* **411**, 195–199 (2001).
5. Haynes, J.D. & Rees, G. Predicting the stream of consciousness from activity in human visual cortex. *Curr. Biol.* **15**, 1301–1307 (2005).
6. Haynes, J.D., Deichmann, R. & Rees, G. Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. *Nature* **438**, 496–499 (2005).
7. Wunderlich, K., Schneider, K.A. & Kastner, S. Neural correlates of binocular rivalry in the human lateral geniculate nucleus. *Nat. Neurosci.* **8**, 1595–1602 (2005).
8. Lee, S.H., Blake, R. & Heeger, D.J. Traveling waves of activity in primary visual cortex during binocular rivalry. *Nat. Neurosci.* **8**, 22–23 (2005).
9. Lee, S.H. & Blake, R. V1 activity is reduced during binocular rivalry. *J. Vis.* **2**, 618–626 (2002).
10. Leopold, D.A. & Logothetis, N.K. Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* **379**, 549–553 (1996).
11. Leopold, D.A., Maier, A., Wilke, M. & Logothetis, N.K. Binocular rivalry and the illusion of monocular vision. in *Binocular Rivalry and Perceptual Ambiguity* (eds. Alais, D. & Blake, R.) (MIT Press, Cambridge, Massachusetts, 2004).
12. Gail, A., Brinksmeyer, H.J. & Eckhorn, R. Perception-related modulations of local field potential power and coherence in primary visual cortex of awake monkey during binocular rivalry. *Cereb. Cortex* **14**, 300–313 (2004).
13. Wilke, M., Logothetis, N.K. & Leopold, D.A. Local field potential reflects perceptual suppression in monkey visual cortex. *Proc. Natl. Acad. Sci. USA* **103**, 17507–17512 (2006).
14. Tong, F. Primary visual cortex and visual awareness. *Nat. Rev. Neurosci.* **4**, 219–229 (2003).
15. Tong, F., Meng, M. & Blake, R. Neural bases of binocular rivalry. *Trends Cogn. Sci.* **10**, 502–511 (2006).
16. Wilke, M., Logothetis, N.K. & Leopold, D.A. Generalized flash suppression of salient visual targets. *Neuron* **39**, 1043–1052 (2003).
17. Bonneh, Y.S., Cooperman, A. & Sagi, D. Motion-induced blindness in normal observers. *Nature* **411**, 798–801 (2001).
18. Wolfe, J.M. Reversing ocular dominance and suppression in a single flash. *Vision Res.* **24**, 471–478 (1984).
19. Maier, A., Logothetis, N.K. & Leopold, D.A. Context-dependent perceptual modulation of single neurons in primate visual cortex. *Proc. Natl. Acad. Sci. USA* **104**, 5620–5625 (2007).
20. Logothetis, N.K. & Schall, J.D. Neuronal correlates of subjective visual perception. *Science* **245**, 761–763 (1989).
21. Gauthier, C. & Hoge, R.D. BOLD-perfusion coupling during monocular and binocular stimulation. *Int. J. Biomed. Imaging* published online, doi: 10.1155/2008/628718 (2 March 2008).
22. Logothetis, N.K. The neural basis of the blood-oxygen-level–dependent functional magnetic resonance imaging signal. *Phil. Trans. R. Soc. Lond. B* **357**, 1003–1037 (2002).
23. Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* **55**, 131–141 (2007).
24. Iadecola, C. Neurovascular regulation in the normal brain and in Alzheimer's disease. *Nat. Rev. Neurosci.* **5**, 347–360 (2004).
25. Logothetis, N.K. The underpinnings of the BOLD functional magnetic resonance imaging signal. *J. Neurosci.* **23**, 3963–3971 (2003).
26. Niessing, J. *et al.* Hemodynamic signals correlate tightly with synchronized gamma oscillations. *Science* **309**, 948–951 (2005).
27. Viswanathan, A. & Freeman, R.D. Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nat. Neurosci.* **10**, 1308–1312 (2007).
28. 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).
29. Lamme, V.A., Super, H., Landman, R., Roelfsema, P.R. & Spekreijse, H. The role of primary visual cortex (V1) in visual awareness. *Vision Res.* **40**, 1507–1521 (2000).
30. Posner, M.I. & Gilbert, C.D. Attention and primary visual cortex. *Proc. Natl. Acad. Sci. USA* **96**, 2585–2587 (1999).
31. Maier, A., Aura, C. & Leopold, D.A. Laminar differences in perceptual modulation of V1 local field potentials. *Soc. Neurosci. Abstr.* **451.12** (2007).
32. Shmuel, A. & Leopold, D.A. Neuronal correlates of spontaneous fluctuations in fMRI signals in monkey visual cortex: Implications for functional connectivity at rest. *Hum. Brain Mapp.* **29**, 751–761 (2008).
33. Mukamel, R. *et al.* Coupling between neuronal firing, field potentials and FMRI in human auditory cortex. *Science* **309**, 951–954 (2005).
34. Nir, Y. *et al.* Coupling between neuronal firing rate, gamma LFP and BOLD fMRI is related to interneuronal correlations. *Curr. Biol.* **17**, 1275–1285 (2007).
35. Judge, S.J., Richmond, B.J. & Chu, F.C. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res.* **20**, 535–538 (1980).
36. Mitz, A.R. A liquid-delivery device that provides precise reward control for neurophysiological and behavioral experiments. *J. Neurosci. Methods* **148**, 19–25 (2005).
37. Gruetter, R. Automatic, localized *in vivo* adjustment of all first- and second-order shim coils. *Magn. Reson. Med.* **29**, 804–811 (1993).
38. Mansfield, P. Multi-planar image formation using NMR spin echoes. *J. Phys. C Solid State Phys.* **10**, L55–L58 (1977).
39. Cox, R.W. Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* **29**, 162–173 (1996).
40. Leopold, D.A., Murayama, Y. & Logothetis, N.K. Very slow activity fluctuations in monkey visual cortex: implications for functional brain imaging. *Cereb. Cortex* **13**, 422–433 (2003).