

Neurophysiological investigation of the basis of the fMRI signal

Nikos K. Logothetis, Jon Pauls, Mark Augath, Torsten Trinath & Axel Oeltermann

Max Planck Institute for Biological Cybernetics, Spemannstrasse 38, 72076 Tuebingen, Germany

Functional magnetic resonance imaging (fMRI) is widely used to study the operational organization of the human brain, but the exact relationship between the measured fMRI signal and the underlying neural activity is unclear. Here we present simultaneous intracortical recordings of neural signals and fMRI responses. We compared local field potentials (LFPs), single- and multi-unit spiking activity with highly spatio-temporally resolved blood-oxygen-level-dependent (BOLD) fMRI responses from the visual cortex of monkeys. The largest magnitude changes were observed in LFPs, which at recording sites characterized by transient responses were the only signal that significantly correlated with the haemodynamic response. Linear systems analysis on a trial-by-trial basis showed that the impulse response of the neurovascular system is both animal- and site-specific, and that LFPs yield a better estimate of BOLD responses than the multi-unit responses. These findings suggest that the BOLD contrast mechanism reflects the input and intracortical processing of a given area rather than its spiking output.

Since its inception^{1–5} BOLD fMRI has generated interest not only as a tool for mapping brain activation, but also as a means of studying the dynamics of neural networks by tracking fMRI response characteristics across various spatial and temporal scales. However, BOLD responses can only measure haemodynamic changes, such as alterations in flow, blood volume, or intravascular magnetic susceptibility, leaving many unanswered questions concerning the relationship between such cerebral haemodynamic changes and actual neural activation. A number of studies in humans or animals have combined fMRI with electroencephalography (EEG)^{6–9} or optical imaging recordings of intrinsic signals¹⁰. However, optical imaging studies essentially also measure haemodynamic responses¹¹, and thus can offer little direct evidence of the underlying neural activity, whereas EEG studies typically suffer from poor spatial resolution and relatively imprecise localization of the electromagnetic field patterns associated with neural current flow.

Active states of any brain site are characterized by time-varying spatial distributions of action potentials superimposed on relatively slow-varying field potentials. An ideal combination is functional imaging in experimental animals with microelectrode recordings of such potentials, including single-spike responses and field potentials, whereby the latter relate well not only to spike activity but also to subthreshold integrative processes in areas such as dendrites that are otherwise inaccessible. Microelectrode recording methods have been used extensively to obtain a wealth of data about the central nervous system in both anaesthetized and alert animals. Adapting this technique for simultaneous electrophysiological and imaging experiments has the potential to link human studies with a large body of animal research carried out over the past 50 years; such a link may lead to answers that are unlikely to be obtained by using either technique in isolation. We present here the first of such simultaneous recordings, using a new electrophysiological measurement technique as well as new methods to actively compensate for the interference produced by the alternation of the field gradients, which are applied for image generation during echo-planar imaging (EPI) (for a comprehensive description of EPI see ref. 12).

The BOLD signal: reflection of spiking or synaptic activity?

To study the neural origin of the BOLD signal we examined the degree of correlation of the haemodynamic response to single- and multi-unit activity (MUA), as well as to LFP. Because single- and multi-unit responses were typically correlated, we will present largely MUA data in this study. MUA and LFP result from the

dynamic interaction of various synaptic and cellular mechanisms, the former reflecting primarily the output of a neural population (within approximately a couple of hundred microns of the electrode tip^{13,14}), and the latter mostly a weighted average of synchronized dendro-somatic components of the input signals of a neural population (within a few millimetres of the electrode tip^{15,16}). Thus their relationship to BOLD responses not only provides insights into the mechanisms underlying the haemodynamic changes, but it can also help us to better interpret the functional meaning of the activation patterns observed in magnetic resonance imaging.

To elicit visual cortical responses we used polar-transformed chequerboard patterns rotating at 60–180 degrees per s. The direction of rotation was reversed every 1 s to minimize adaptation. The time course of the BOLD signal was studied and correlated with the neural response by carrying out single-slice (2-mm thickness) fMRI using single- or multi-shot EPI with $0.75 \times 0.75 \text{ mm}^2$ in plane actual resolution. Optimal selection of a slice was achieved by determining the electrode position in an anatomical scan and then running a functional multi-slice (13 slices, 2-mm thickness), multi-shot EPI scan using a block design stimulation model. A new recording method, specially constructed electrodes and micro-manipulators, and a system to compensate for interference were developed to measure neural activity in the presence of the strong, alternating gradients of the magnetic field (see Methods).

Figures 1a and b show the location of the electrode tip and the BOLD response elicited by the visual stimulus in the primary visual cortex of one animal (B00). Figure 1c shows the comprehensive broadband neural signal (that is, the de-noised signal of 50 mHz to 3 kHz bandwidth, which includes LFP, MUA and single spikes) superimposed on the BOLD activation for the region of interest (ROI) around the electrode tip, that is, the average time series of the voxels inside the dotted green line. The average ROI used in the experiments was $7.1 \times 2.3 \text{ mm}^2$ (1 s.d. = 3.5 mm^2), and was defined as a small region around the tip of the electrode containing only the primary visual cortex. For any given voxel of this region the intensity of the BOLD image comprises a time series of data. Time series within this region that significantly correlated with the time course of the stimulus ($r > 0.35$; $P < 0.00001$) were selected for further analysis. The neural signal typically showed an increase of overall activity immediately on presentation of the visual stimulus that lasted for the entire duration of stimulation. The yellow trace shows the root mean square (r.m.s.) value of successive 250-ms windows.

The window size was chosen to equal the repetition time of the pulse sequence. The r.m.s. typically showed a transient, 2–3-s-long increase followed by a sustained response and then a slight decrease of activity that was also noticeable in the baseline of the de-noised raw signal. The mean onset (the time point at which the BOLD signal first increased above 3 s.d. of the baseline) of the BOLD signal was 2.16 s (1 s.d. = 0.94 s) after neural activation, reaching a plateau about 7 s later. To examine the contribution of MUA and LFP signals to the haemodynamic response we applied time-dependent frequency analysis to the raw data. Figure 1d–f shows the magnitude of the time-dependent Fourier transform (spectrogram) of the signal computed using a sliding window (Hamming window of 250 ms) for three different stimulus-presentation times (24, 12 and 4 s). Typically, after stimulus presentation a transient increase in power was observed across all frequencies, followed by a lower level of activation that was maintained during the entire stimulus presentation. A prominent characteristic in all spectrograms was a marked stimulus-induced increase in the magnitude of the LFP, which was always larger than that observed for MUA. The vertical panel across the time axis of each three-dimensional plot shows the average LFP and MUA responses computed as the mean vector of the time series between 40–130 Hz (black dashed lines) and 300–1,500 Hz (blue dashed lines) frequency regions. The vertical panel along the frequency axis shows the average spectra for the pre-stimulus, stimulation, and post-stimulus periods. A decrease in neural activity was also observed immediately after the termination of the stimulus, as shown by the deep-blue regions in all spectrograms.

To confirm these results across all scans, a spectrogram (Fig. 2) was computed for the first 6 s of the neural response averaged over all data collected during 24, 12.5, 12 and 6-s stimulus presentation.

The time series at each frequency is expressed in units of the standard deviation (s.d.) of the activity in the pre-stimulus period, and it therefore represents the signal-to-noise ratio of the response at that frequency. The maximum increase in power over all sessions (10 monkeys, $N = 619$ experiments) was found at 72.96 Hz (1 s.d. = 21.04 Hz), in other words within the gamma range of the LFPs. The top panel shows the mean LFP, MUA, total (over all frequencies) and BOLD responses for all data. There is a marked difference in signal-to-noise ratio (scale of the vertical axes) between the neural and the BOLD signals. Such a difference can, in principle, result in statistical rejection of the activation of various regions during mapping experiments, despite the fact that the underlying neural activity is highly robust and significant.

Approximately one-quarter of all multiple-unit responses were found to be transient, in that they showed an initial increase in amplitude, returning to baseline within 2–4 s. Figure 3a shows simultaneous recordings of haemodynamic responses, LFPs and transient single- and multi-unit activity in striate cortex. Both the spike-density function, representing a neuron’s instantaneous firing rate, and the MUA show strong adaptation, returning to the baseline around 2.5 s after stimulus onset. In contrast, the activity underlying the LFPs remains elevated for the duration of the visual stimulus. There was no single observation period or recording site for which the opposite result was observed, namely a highly correlated MUA signal and an uncorrelated or missing LFP signal. Similarly at no time did we observe MUA that was larger in magnitude or signal-to-noise ratio than the measured LFP activity. These findings suggest that BOLD activation may actually reflect more the neural activity related to the input and the local processing in any given area, rather than the spiking activity commonly thought of as the output of the area.

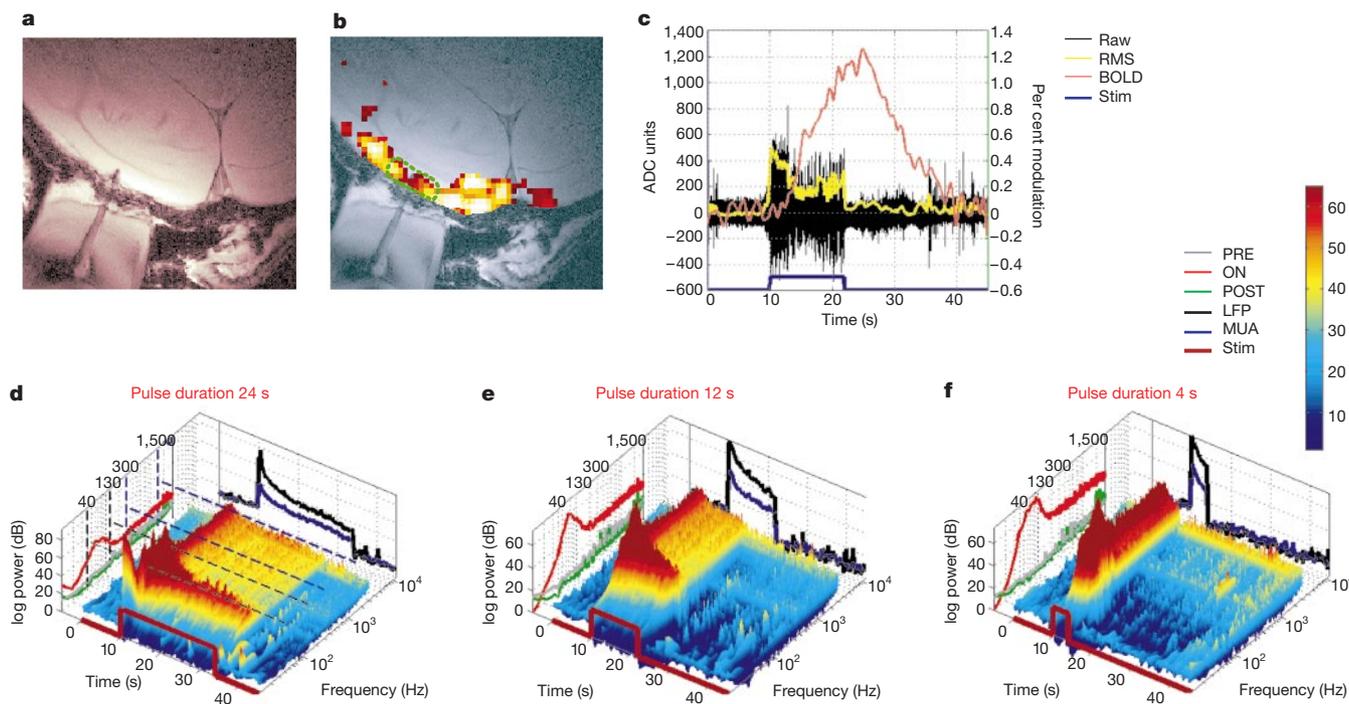


Figure 1 Neural and BOLD responses to pulse stimuli. **a**, FLASH scan (see Methods) showing the location of the electrode tip in primary visual cortex. **b**, BOLD response to rotating checkerboard patterns in striate cortex. Activation can be measured around the electrode tip. **c**, Haemodynamic response (red) superimposed on the de-noised raw neural signal (black). The term ‘de-noised raw’ denotes that no other signal processing beyond the removal of gradient interference (see Methods) was done. The r.m.s. of the signal is indicated by a thick yellow line. **d–f**, Spectrograms for data collected over 24, 12

and 4 s. In each three-dimensional plot, the vertical panel along the time axis shows the average LFP and MUA responses, namely the mean vector of the time series between black and blue dashed lines, respectively. The vertical panel along the frequency axis shows the average spectra for the pre-stimulus, stimulation, and post-stimulus periods. Colour bar shows the logarithm of power. ADC, Analogue to digital converter; STIM, time course of the visual stimulus; PRE, pre-stimulus period; ON, stimulus presentation period; POST, post-stimulus period.

The stronger contribution of LFPs in BOLD responses might be the result of differences in spatial summation, as LFPs usually integrate signals from a couple of millimetres, whereas MUAs do so only for a few hundred micrometres. To test whether or not such an explanation is plausible we repeated the same experiments outside the magnet, with exactly the same stimulation and anaesthesia conditions, but with a multi-unit recording system¹⁷. Intracortical recordings were carried out with a 4 × 4 array of microfibre electrodes (quartz Pt₉₀W₁₀; 80 μm shaft diameter, spacing of 250 μm centre to centre, impedance 250–750 kΩ, at 500 Hz). MUAs of distant electrodes were added using a weighting factor that decreases with the same rate as the LFPs decrease with distance from the electrode¹⁸. We (arbitrarily) assumed that MUAs are collected over an area of 200 × 200 μm², and that LFPs are collected over an area of 2,000 × 2,000 μm², which means that about 100 MUA signals (multiple recordings from each electrode) must be ‘summed’ to account for the spatial summation assumed for LFP. The contribution of MUA may even decrease by summing this type of activity. This is presumably due to the lack of significant synchronization in this frequency band. Summing the signals of a single trial slightly increased the transient portion of the MUA, but it never increased the sustained part of the response clearly seen in the LFP band (see Supplementary Information).

Finally, the transfer function of the entire recording system was carefully examined to ensure that lower frequencies were not preferentially amplified due to impedance differences at different frequencies. Both measurements and simulations showed that the gain of a frequency component monotonically increases as the frequency increases. That is, the transformation of the linear

transfer function into a constant function would attenuate MUA even more with respect to the LFP signals (see Supplementary Information).

Modelling the BOLD response from neural and imaging data

The simultaneous recording of the neural and haemodynamic responses enable the application of system-identification techniques on a trial-by-trial basis to build a model of the BOLD mechanism. We made an initial attempt to describe this mechanism by estimating its impulse response to the underlying neural activity. Correlation analysis was applied to both the measurements conducted during visual stimulation and the measurements of spontaneous activity. In either case, the input (neural) data was pre-whitened to make the results as uncorrelated as possible. Pre-whitening was done by fitting a 10th-order autoregressive model to the neural data (it corresponds to filtering ‘deviations (peaks) from whiteness’ with a filter of 60 dB oct⁻¹ attenuation and zero phase shift). The parameters of the autoregressive model were estimated by using the least-square method (that is, by minimizing the standard sum of squared, forward prediction errors). Correlation analysis was applied after subjecting the output (haemodynamic) data to the same filter. The impulse response is the cross covariance function (positive lags) of the neural and BOLD responses. Figure 4a shows the impulse functions computed from the LFP and BOLD responses from the striate cortex of two monkeys (b97, k00) under conditions of no stimulation. They were used to convolve the LFP responses to a stimulus of 100% contrast presented for 12 s (Fig. 4b) and to a stimulus with the contrast changing every 4 s (Fig. 4d). In a first approximation the model predicts the measured response quite well (session b972y1: $r^2 = 0.905$ with LFP, $r^2 = 0.725$ with MUA; session k006i1: $r^2 = 0.769$ with LFP, $r^2 = 0.625$ with MUA). However, the determination coefficient (normalized residuals) decreased rapidly with increasing stimulus duration, indicating the existence of higher-order terms ignored by linear systems analysis. Figure 4c plots the distribution of determination coefficients for BOLD estimates from LFP and MUA, using impulse

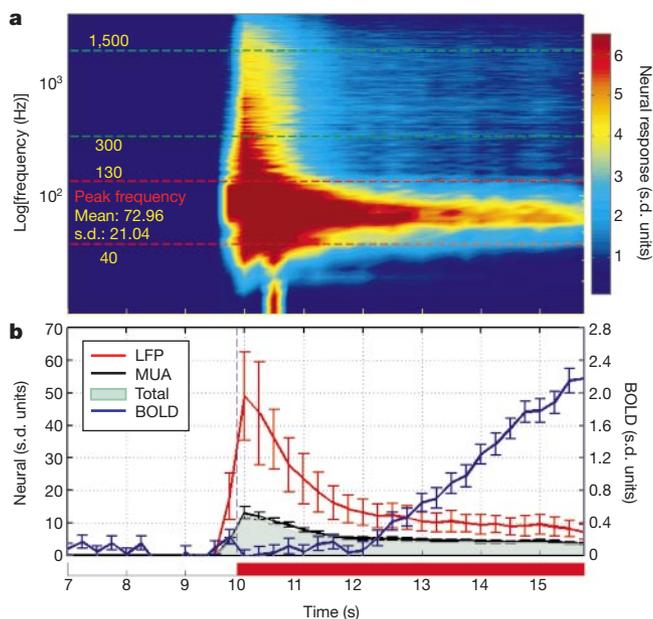


Figure 2 Time-dependent frequency analysis for population data. **a**, Spectrogram of the first 6 s of the neural response averaged over all data collected during 24, 12.5, 12 and 6 s of stimulus presentation (10 monkeys, 619 experiments). Each time course is expressed in units of the standard deviation of the pre-stimulus period. Colour (and thus the colour bar) encodes the reliability of signal change for each frequency (signal-to-noise ratio) rather than the magnitude of the power spectrum. Red and black dashed lines show the LFP and MUA frequency bands, respectively. **b**, Mean LFP (red), MUA (black) and total (green surface) neural response (average across all frequencies), together with the BOLD signal (blue). The total signal is very similar to the MUA signal as the LFP represents a small frequency range (note the logarithmic scale of the frequency axis in **a**). The figure shows the significantly higher LFP activation for both the transient and the sustained portion of the response. Error bars are 1 s.d. with $N = 10$ (number of monkeys).

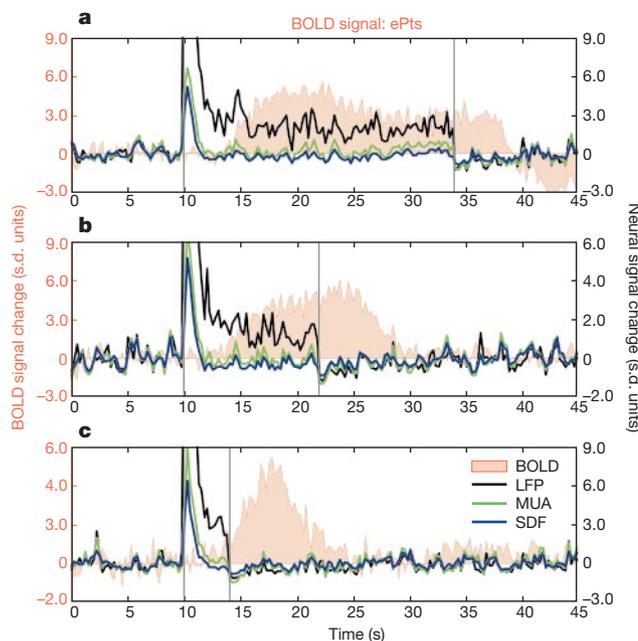


Figure 3 Simultaneous neural and haemodynamic recordings from a cortical site showing transient neural response. **a–c**, Responses to a pulse stimulus of 24, 12 and 4 s. Both single- and multi-unit responses adapt a couple of seconds after stimulus onset, with LFP remaining the only signal correlated with the BOLD response. SDF, spike-density function (see text); ePts, electrode ROI—positive time series.

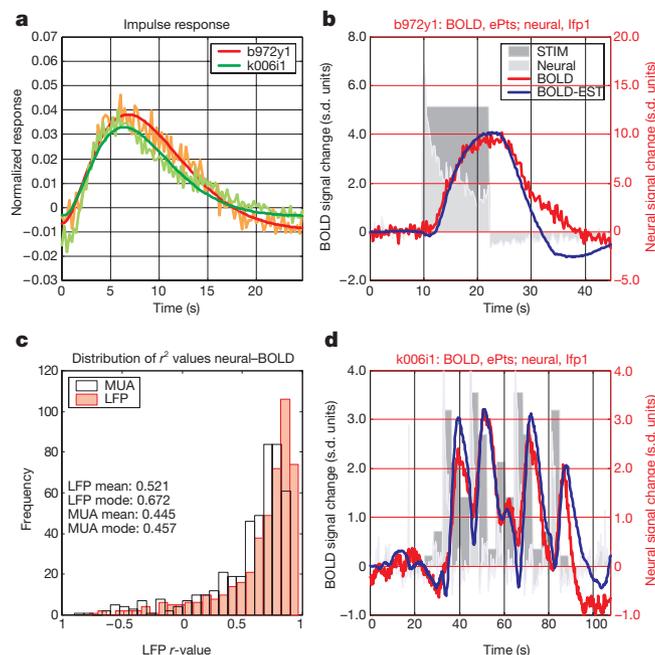


Figure 4 Correlation analysis for the estimation of the impulse response of the neurovascular system and validation of data collected with a pulse or a variable-contrast stimulus. **a**, Impulse responses for two sessions. Each computed by correlating the LFP and BOLD signal. **b, d**, Model validation. The red curve shows the measured and the blue the estimated BOLD response. The dark grey rectangle(s) show the contrast level at any given time, and the light grey curves show the neural response. **c**, Distribution of the

coefficient of determination (r^2) of LFP and MUA. The r^2 values for LFP were significantly higher than for MUA. Residual analysis (data not shown) similarly showed lower errors (in the least-square sense) in LFP-based estimates than estimates made on the basis of MUA responses. The plot in **c** includes data collected with pulse stimuli of 24, 12, 6 and 4 s ($N = 460$). ePts, electrode ROI—positive time series; lfp1, LFP from 40 to 130 Hz (see red dashed lines in Fig. 2a).

functions computed from the LFP/BOLD and MUA/BOLD covariance, respectively. On average, LFP response gave a better estimate of BOLD responses (in terms of least-square error or r^2 value) than did MUA (r_{LFP}^2 mean = 0.521, mode = 0.672; r_{MUA}^2 mean = 0.445, mode = 0.476; $P < 0.00016$, t -test).

Neural and BOLD responses as a function of stimulus contrast

To elucidate the relationship between BOLD responses and the underlying neural activity, we used four stimulus-contrast levels with four different stimulus-presentation durations to estimate the impulse response of the BOLD signal. Figure 5 shows data for a stimulus duration of 12.5 s; Fig. 5a shows fMRI responses to this stimulus. Each curve is the mean time course of the BOLD response averaged across all voxels in a ROI around the electrode tip and across all repetitions ($N = 4$). To compute the amplitude of the response a nonlinear function ($S \times \text{gamma function} \times \text{cosine}$) was fitted to the data (smooth lines in Fig. 5a). The scale S of this function was taken to be the response amplitude. Amplitudes are given as a fraction of the maximum response. As shown previously, the fMRI responses increased with stimulus contrast; this increase is directly attributable to analogously increased neural activity (Fig. 5b). Both LFP and MUA (data not shown) increased with stimulus contrast, although at a different rate than the haemodynamic responses. The thick lines in Fig. 5b were obtained by convolving the neural responses with the corresponding impulse-response function. The maximum value of these curves was taken as the neural response amplitude; it is plotted against contrast in Fig. 5c together with the BOLD responses. Neither neural nor fMRI response is a linear function of contrast, but both increase monotonically with stimulus contrast. A ‘threshold’ type of nonlinearity is evident for the low-contrast values, in that the fMRI response reaches about 50% of its maximum amplitude even with the lowest contrast tested (12.5%). The normalized BOLD response as a function of LFP and MUA are plotted in Fig. 5d. With the tested

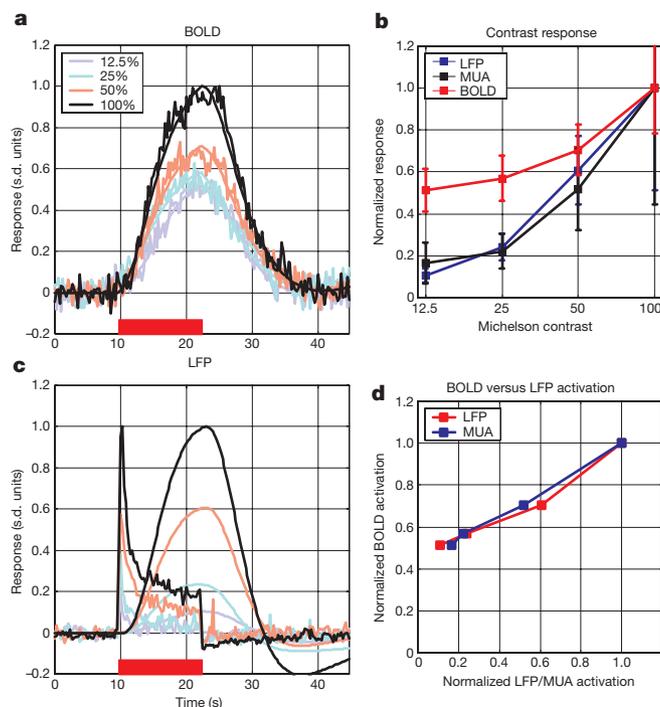


Figure 5 fMRI responses to pulse stimuli at four different contrasts (12.5, 25, 50 and 100%). **a**, Mean fMRI response superimposed with a model estimated with nonlinear curve fit. The scale parameter of the model was taken as the response amplitude. **b**, Normalized response amplitude of LFP and MUA against contrast. Data from five sessions with a pulse duration of 12.5 s. **c**, LFP responses for four different contrasts. Smooth lines are the result of convolution of the neural responses with the impulse response estimated by correlation analysis. **d**, Normalized BOLD response as a function of LFP and MUA. Responses were normalized by dividing each response by the maximum response.

contrasts, the relationship between the BOLD and the neural responses was found to be roughly linear.

Discussion

We have examined the relationship between the BOLD fMRI signal and the underlying neural activity in simultaneous intracortical electrophysiology and imaging experiments in anaesthetized monkeys. Our results show unequivocally that a spatially localized increase in the BOLD contrast directly and monotonically reflects an increase in neural activity. The time course of the haemodynamic response is roughly a low-pass-filtered expression of the total neural activity, given here as r.m.s. values in each 250-ms window of the frequency- and amplitude-modulated neural signal. Earlier studies have described the time course of the BOLD signal obtained during brief visual stimulation. After stimulus onset, the BOLD signal in humans, rats, cats or monkeys often exhibits an initial dip^{19,20}, which is attributed to the rapid increase in local deoxyhaemoglobin observed in optical imaging studies²¹. The subsequent signal increase is delayed by 2–3 s, followed by a ramp of 6–12 s to a plateau or peak value for long (>10 s) or short (<10 s) pulses, respectively. The signal returns to the baseline with a similar ramp, although often a prolonged post-stimulus undershoot is evident, which is attributed to volume changes^{20,22–24}. Our data (not presented here) show that this undershoot is often preceded by a marked inhibition in the neural response, so it may be related, at least in part, to changes in neural activity.

To better understand the neural mechanisms underlying the BOLD response, spiking and synaptic activity were examined separately by analysing single- and multi-unit activity and LFPs, respectively. LFPs are often dominated by stimulus-induced and usually stimulus-locked fast oscillations in the range of 30–150 Hz^{25–28}, as are human EEGs or magneto-encephalograms during visual or auditory tasks^{29,30}, and their role in sensory and perceptual information processing and multimodal sensory or sensory-motor integration has been debated extensively^{31,32}. Our results show that the increase in LFPs during stimulation is significantly stronger than that of multi-unit activity. Moreover, while MUA was often found to adapt, returning almost to baseline levels, LFP activity was always maintained throughout the stimulus presentation. The fast MUA decay observed in about one-quarter of the recording sites may be, at in least in part, due to lack of directional selectivity in those sites, which results in adaptation regardless of the changes in rotation direction of the stimulus. The average LFP response was also found to give better estimates of the BOLD signal (smaller error in the least-square sense) when neural activity was convolved with the neural-vascular impulse response function than when MUA was used as the system's input. In all of the measurements, the signal-to-noise ratio of the neural signal was an average of at least one order of magnitude higher than that of the fMRI signals. This observation indicates that the statistical analysis and thresholding methods applied to the haemodynamic responses probably underestimate a great deal of actual neural activity related to the stimulus or task, and suggest that a certain degree of caution is called for when interpreting mapping studies, particularly when precise localization of activity is required.

The greater contribution of LFP activity to the fMRI signal is consistent with findings regarding the bioenergetics underlying this signal. It is well established that neural activity and energy metabolism are tightly coupled³³. A quantitative relationship can actually be established between imaging signals and the cycling of certain cerebral neurotransmitters^{34–36}, as synaptic activity is tightly coupled to glucose uptake^{37,38}. NMR spectroscopy experiments exploiting such a coupling showed that the energy demands of glutamatergic neurons account for 80–90% of total cortical glucose usage in rats³⁹ and humans⁴⁰. The present findings also imply that the greater portion of the haemodynamic signal changes reflect the energetically expensive synaptic activity such as that related to the

LFP signals. Both our physiological measurements and the spectroscopy results are incompatible with models suggesting a quantitative relationship between the spike rate of neurons and the haemodynamic response⁴¹.

A number of studies have explored the contrast-response function of single cells and the relationship of the fMRI signal to the contrast. The responses of single units vary from cell to cell. In general, neurons show an increasing response that is nonlinear at low contrasts, and a linear increase up to a certain contrast level at which the response reaches its asymptote^{42,43}. Like the neural responses, fMRI BOLD response was also found to be a nonlinear function of stimulus contrast^{44,45}; however, a linear systems analysis on the fMRI responses predicted a linear relationship between the BOLD and neural activity⁴⁵. The data presented here are consistent with these studies. Furthermore they directly show that a linear relationship between the magnitude of the neural and the BOLD signals exists at the contrasts tested (12.5–100%). Nevertheless, although increases in stimulus contrast lead to monotonically increasing BOLD levels in primary visual cortex, manipulation of stimulus detectability using noise can lead to non-monotonic BOLD tuning⁴⁶.

The BOLD contrast mechanism directly reflects the neural responses elicited by a stimulus. In a first approximation BOLD and neural responses are shown to have a linear relationship for short stimulus presentation durations. Neural signals are characterized by a considerably higher signal-to-noise ratio than the haemodynamic response, suggesting that the extent of activation in human fMRI experiments is very often underestimated to a significant extent owing to the variation in the vascular response. Finally, the haemodynamic response seems to be better correlated with the LFPs, implying that activation in an area is often likely to reflect the incoming input and the local processing in a given area rather than the spiking activity. Although it is reasonable to expect that output activity will usually correlate with neurotransmitter release and pre- and post-synaptic currents, when input into a particular area primarily has a modulatory role, fMRI experiments may reveal activation in areas in which no single-unit activity is found in physiological experiments. □

Methods

Surgery, anaesthesia for fMRI and visual stimulation

This study involved 29 combined, electrophysiological fMRI experimental sessions in 10 healthy, anaesthetized monkeys (*Macaca mulatta*, 6–9 kg). All studies were carried out with great care to ensure the well being of the monkeys; were approved by the local authorities (Regierungspraesidium); and were in full compliance with the guidelines of the European Community for the care and use of laboratory animals. A detailed description of the surgical procedures has been published elsewhere²⁰ (see also Supplementary Information).

MRI data collection

We made measurements on a vertical 4.7-T scanner with a 40-cm diameter bore (Biospec 47/40v; Bruker). The system was equipped with a 50 mTm⁻¹ (180 μs rise time) gradient coil that was actively shielded (Bruker, B-GA 26), with an inner diameter of 26 cm. A primate chair and a special transport system were designed and built for positioning the monkey within the magnet. Customized, small radio-frequency coils (30–80 mm diameter) were used (Fig. 6a), which were optimized for increased sensitivity over a given ROI, such as a portion of the primary visual cortex or regions of areas V2, V3, V4 and V5 (middle temporal visual area). The coils were attached around the recording chamber and were used as transceivers. All images were acquired using a 96 × 96 mm field of view. T1-weighted, high resolution (256 × 256 matrix, 0.5 mm thickness) anatomical scans were obtained using the three-dimensional, modified driven equilibrium Fourier transform⁴⁷ (MDEFT) pulse sequence, with an echo time of 4 ms, repetition time of 14.9 ms, flip angle of 20°, and 4 segments. To position the electrodes, 45-slice anatomical scans were acquired with the fast, low-angle shot (FLASH) sequence⁴⁸ with an echo time of 8.9 ms; repetition time of 2,000 ms; field of view of 96 × 96 mm; and a 512 × 384 matrix, reconstructed to 512 × 512, slice thickness of 0.5 mm.

Multi-slice fMRI was carried out by using multi-shot (segmented) gradient-recalled EPI (GE-EPI)⁴⁹. To examine the distribution of activation around the electrode, we initially collected volumes of 13 slices (2 mm thickness), each with a field of view of 96 × 96 mm on a 128 × 128 matrix. The acquisition parameters were: echo time of 20 ms; repetition time of 750 ms; flip angle of 40°; EPI zero phase of 8.192 ms or 40% of phase steps; pulse length

of 3.0 ms; spectral width of 100 kHz; line acquisition time of 1.28 ms; 8 segments; segment acquisition time (MRI readout window width) of 20.48 ms; repetition time between slices of 50.36 ms; and one excitation per phase-encoded step.

To study the time course of the BOLD signal and directly compare it to the neural responses, fMRI was carried out with GE-EPI with 1 or 8 interleaved segments, each segment acquired during a separate visual stimulation epoch using a 128×128 matrix, echo time/repetition time of 20/250 ms; flip angle of $20\text{--}25^\circ$; pulse length of 3.0 ms; spectral width of 100–150 kHz; EPI zero phase at 20% of phase steps; readout window width of 20.48 ms; and one excitation per phase-encoded step. The segments were

acquired separately²⁰ and merged into one image by sorting all K-space (frequency-phase-space) lines according to their phase. In these segmented acquisitions, echo time shift was used to minimize ghosting (image artefacts due to imperfections in the gradient pulses or gradient-induced eddy currents). The time of echo shift for a segment with the number j was calculated as $(1+j) \times \text{aqt} / (\text{number of EPI segments})$, where $\text{aqt} = \text{matrix size} / (2 \times \text{spectral width})$ (aqt is the readout time for a single echo). To minimize the effects of inflow and of large drainage vessels we consistently used flip angles that were smaller than the computed Ernst angle by 10° . In each session and for each monkey, an autoshim algorithm was used to tune the linear shim coils, followed by the manual tuning of the higher-order shim coils. Shimming (optimization of the homogeneity of the main magnetic field) was typically performed with a hard pulse of $50 \mu\text{s}$ duration and a receiver bandwidth of 50 kHz. Depending on the radiofrequency coil, either global or local volume shimming (FASTMAP⁵⁰) was used. The line width (full width at half height) at the end of the shimming procedure ranged from 30 to 70 Hz for the entire scanned volume. Together with the image acquisition both the plethysmogram (detection of volume changes due to arterial pulsations, roughly corresponding to an ECG signal) and the flow signal of the ventilator were digitized at 250 Hz, and stored for later evaluation and removal of physiological artifacts from the BOLD time courses.

MRI data analysis

All data were analysed off line on Pentium computers running the Linux or NT operating system. Analysis was performed using our own software (MATLAB) and the MEDx 3.0 image-processing package (Sensor Systems). After normalization of the data, the linear trends were removed and the data were wavelet filtered to remove temporal noise. Respiratory and cardiovascular artifacts were removed by examining the spectra of the recorded plethysmogram and respiratory signal, and creating appropriate filters. Respiration was kept constant in all experiments to 24 strokes per min, corresponding to 0.417 Hz (repetition time of 250 ms; sampling rate was 4 Hz). To increase statistical significance, all images used for three-dimensional brain reconstruction (voxel size of $1 \times 1 \times 2 \text{ mm}^3$) were spatially filtered (convolution filter with full width at half height of 1.5 mm, on a 3×3 kernel—roughly 95% of the filter is in the kernel). Both intergroup statistics (parametric paired t -test or F ratio) and correlation analysis were used to generate functional maps.

Electrophysiological recording

We developed all of the recording hardware, including electrodes, microdrives, signal conditioning and interference compensation devices (N.K.L., A.O. and M.A., manuscript in preparation). A plastic chamber (see Supplementary Information) formed to fit precisely the skull of a monkey was implanted over the occipital pole during an aseptic surgical procedure. The skull inside the chamber was left intact. At the beginning of each experiment a 2 mm trephination was performed through which the electrode was introduced into the brain. Electrodes were made of platinum-iridium ($\text{Pt}_{90}\text{Ir}_{10}$) wire etched with sodium cyanide (NaCN) solution and coated with glass (Corning glass 7570). The glass-coated tip was glued into a glass capillary tube (1.5 mm outer diameter) with super glue. The wire extended 5–10 mm beyond the end of the capillary tube and served as the contact point for the electrode. The electrode holder (Fig. 6a) consisted of three concentric, metallic cylinders (copper beryllium or a bronze alloy, CuSn_7PbZn). The innermost cylinder served as the contact point for the electrode, the middle as far-interference sensor, and the outermost layer as the amplifier ground (see below). We insulated the cylinders from each other with layers of polyetheretherketone (PEEK). The concentric cylinders, in particular the outermost cylinder, offer the additional advantage of being a rotation-symmetrical shield for the electrode, permitting optimal ground contact to the animal while avoiding any kind of loops susceptible to induction. We sealed the gap between electrode and electrode holder with silicon gel. Electrode, sensor and ground were connected to the amplifier with two twisted, coaxial cables. The outer conductor of both cables was connected to the ground cylinder of the electrode holder. The electrode holder was held firmly in a microdrive attached to the recording chamber by means of a holding ring. The positioning of the electrode was accomplished by manually turning a screw with a thread pitch of $700 \mu\text{m}$. In addition, a three-coil, magnetic-field sensor (see below and Fig. 6a) was mounted on the microdrive. The joints between microdrive, holding ring and chamber were sealed with silicon gel. The chamber was filled with deuterium saline (0.9% NaCl in D_2O) to provide electrical contact between the ground cylinder and the animal. Deuterium saline was used instead of regular saline to avoid changing the Q-factor of the radio-frequency coil and to permit optimal shimming. Furthermore, 0.6% agar was added to the deuterium saline to minimize oscillations in both the saline and the dura during the alterations of the readout gradient. At the beginning of each session the electrode was lowered into visual cortex and positioned such that both signal intensity and stability were maximized. On the basis of magnetic resonance images most recordings were obtained from granular and infragranular layers of V1.

Signal acquisition and conditioning

Several problems arise when one tries to apply conventional neurophysiology techniques in the context of magnetic resonance imaging. Typically, a pre-amplifier is positioned near the head of the animal and is connected by means of a cable to a main amplifier some distance away. However, the strong magnetic field alternations required for the magnetic resonance image formation incapacitate the pre-amplifier by inducing voltages in any existing loop in the circuit. Shielding reduces electrical interference, but it is ineffective against magnetic interference, whereas materials such as Mu-metals that attenuate magnetic interference contain iron, which affects field homogeneity. Furthermore, low-noise, low-frequency voltage amplifiers typically used in neurophysiology rectify high-frequency voltage signals coupled with the input signal, an effect that occurs at all stages of

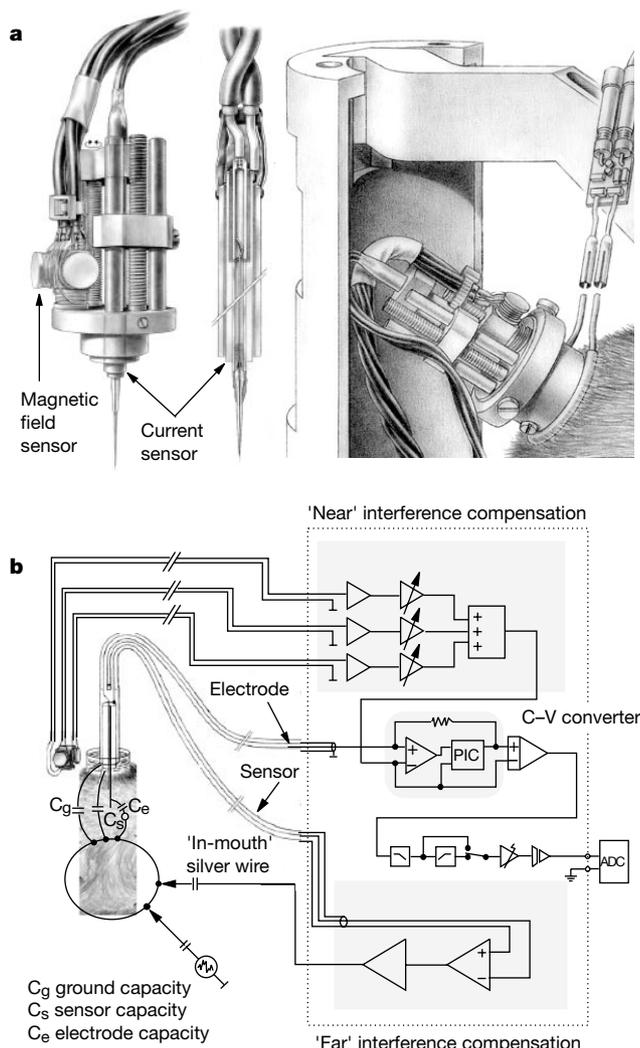


Figure 6 Recording hardware. **a**, Recordings were conducted with glass-coated, platinum-iridium electrodes. Their holder consisted of three concentric, metallic cylinders, the innermost serving as the contact point for the electrode, the middle as the far-interference sensor, and the outermost as the amplifier ground (see Methods). The cylinders were insulated from each other by layers of PEEK. Electrode, sensor and ground were connected to the amplifier with two twisted, coaxial cables (middle panel). The electrode was positioned by manually turning a screw with a fine thread pitch. In addition, a three-coil magnetic-field sensor was mounted on the microdrive to compensate for near interference (see Methods). The radio-frequency coil used for imaging was placed around the chamber. **b**, Block diagram of the recording and compensation circuitry. The animal can be conceived as being capacitively connected to any metal contact, including connections to ground (C_g), sensor (C_s) and electrode (C_e). Because of the finite animal-to-ground capacity, a fraction of the interference currents flow through the electrode. To compensate for such currents we used a sensor built into the electrode holder (current sensor). Interference originating near the electrode tip or within the electrode holder and the cables was compensated for by using three small, orthogonally oriented coils that were identical and were positioned near the electrode (magnetic field sensor). C–V, current to voltage converter; ADC, analogue to digital converter; PIC, proportional integral controller.

integrated circuits. In a 4.7-T magnet, amplifiers immediately rectify the 200 MHz radio-frequency excitation pulses, resulting in substantial disturbance of the output that only recovers several tens of milliseconds after the radiofrequency signal has subsided. Moving the conventional pre-amplifier outside of the gradient tube does not solve the problem, as the capacitance of the cable between the electrode and the pre-amplifier—in our case a cable of at least 2 m—together with the electrode impedance, act as a voltage divider compromising the signal.

To avoid signal loss owing to increased cable length, we developed a method of measuring current instead of voltage that is insensitive to cable length. Specifically, current flow from the electrode tip was measured over an ultimate cable length of 6 m and converted into voltage at the input stage of the amplifier, with pre-amplifiers, power amplifiers, operating elements and interference compensation circuits all packed into one device located outside the magnetic field. Details of the circuitry will be described elsewhere (N.K.L., A.O. and M.A., manuscript in preparation). Under our measurement conditions (roughly 600 pF input capacitance owing to cable length and high impedance of the feedback circuit of the pre-amplifier) common current-to-voltage converters are unstable. Instead of a classical current-to-voltage converter, a voltage amplifier, a proportional-integral controller (PIC) and a 10 MΩ resistor were used to ensure stability under the desired bandwidth conditions (Fig. 6b). The PIC regulates its output voltage so that the input voltage of the amplifier ultimately equals 0 V. As a result the output voltage equals the voltage drop across the 10 MΩ resistor and is therefore proportional to the measured input current. Furthermore, no signal loss occurs by charging the cable capacity. To obtain the fastest possible impulse response with the minimum possible overshoot for the desired cable length, the tuning parameters (amplification and cut-off frequency) of the Proportional Integral Controller were adjusted with a square-wave current.

Interference sources and far interference compensation

Two types of interference had to be compensated for: interference originating from a distance greater than that from the electrode tip to the electrode ground (far interference), and interference originating from the immediate vicinity of the electrode tip (near interference).

Owing to the metal–electrolyte interface the animal is capacitively connected to any metal contact (Fig. 6b), including connections to ground and electrode. Because the animal-to-ground capacity is finite, a fraction of the interference currents flowing towards the animal (for example, currents through ECG or intravenous lines) will also flow through the electrode. To compensate for such currents we used a sensor built into the electrode holder. The current between the sensor and the ground was measured, amplified and converted into a voltage. The voltage was inverted, amplified again in two steps—integrated output current limitation for the protection of the animal (10 mA maximum, 2 V maximum)—and was sent back to the animal through a junction electrode in the cavity of the mouth. This compensation set the voltage of the mouth electrode so that the current flowing through the sensor electrode remained at 0 A. In other words, the mouth current equalled the sign-inverted interference current flowing to the animal. As the three capacitances (Fig. 6b) are parallel, setting the sensor current to zero forces all the other currents to zero as well, thereby eliminating the interference measured at the tip of the electrode, and leaving the neural signal unaffected.

Near interference compensation

Interference originating and acting near the electrode tip or originating within the electrode holder and the cables cannot be detected by the sensor and compensated for by the circuit described above. The magnitude of this interference is sufficiently large to require a dedicated compensation circuit. Therefore, local changes in the magnetic field caused by the alternating field gradients during image acquisition were monitored by three small, orthogonally oriented coils that were identical and were positioned near the electrode, and the measured signal was added to the ‘ground’ of the current-to-voltage converter to neutralize the interference. This method of compensation is effective against induction currents in the electrode holder and the cable, which could not be completely avoided by the rotation-symmetrical construction described above. Furthermore this method is effective against induction currents owing to the distance between the electrode tip and ground (possibly because of eddy currents in the saline in the recording chamber and in the brain). Briefly the operation principle of this circuit is as follows. The voltages induced in the coils serve as the inputs to the compensation circuit. These are amplified and passed through an additional amplification stage, where the gain can be adjusted with a precision potentiometer to a value between –1 and 1. The sum of the three independently adjusted signals serves as the reference (ground) for the current-to-voltage converter. The orthogonal orientation of the small coils and the adjustability of sign and value of the gain make it possible to simulate the induction voltage in a wire loop of any given diameter and orientation (vector addition). A virtual wire loop that has opposite sense of winding can be adjusted by the precision potentiometers in such a way that the remaining loops, caused by asymmetries in electrode holder and cable, are effectively compensated.

Neural data analysis

Signals were amplified by 3–30 mV pA⁻¹. In a conventional voltage-measuring system using an electrode of 300 kΩ impedance measured at 1 kHz, this would amount to a total amplification of 10³ to 10⁵. The bandwidth of the main amplifier was 50 mHz to 3 kHz (12 dB oct⁻¹ and 18 dB oct⁻¹, respectively) and the signal was digitized with 22.3 kHz using a 16-bit AD converter, set to ± 10 V input range. It was subsequently decimated by a factor of three to 7.43 kHz. From the broadband recordings three signals were extracted for analysis: (1) single-unit trains of action potentials convolved with a gaussian of fixed kernel (s.d. = 5 ms) to yield spike-density functions; (2) MUA in the range of 300–3,000 Hz; and

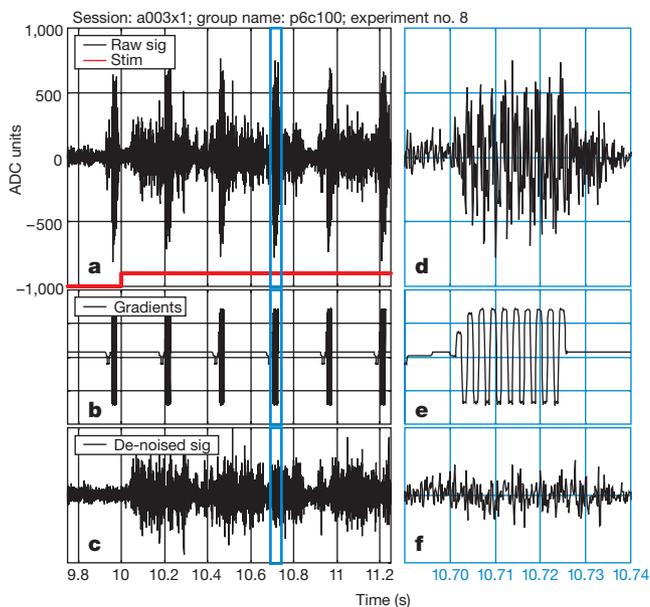


Figure 7 Elimination of residual interference by applying PCA (see Methods). **a–c**, A 1.4-s-long segment of the neural signal; **d–f**, magnification of the signal part included within the blue lines. The upper panels show the raw data, middle panels the recorded gradient currents, and lower ones the de-noised signal. In **a** and **d**, cell activity is shown superimposed on the strong interference induced by the gradient coils. The periodic alternations in **b** and **e** are due to the switching of the readout gradient.

(3) LFPs in the range of 10–130 Hz. Band separation was accomplished either by filtering (10–130 Hz, 36 dB oct⁻¹ and 300–3,000 Hz, 48 dB oct⁻¹, both with zero phase shift) or by selectively averaging the spectrogram (generated with a Hamming window of 250 ms) of the signal across frequency. In the first case, MUA was rectified and low-pass filtered (150 Hz, 24 dB oct⁻¹, zero phase shift) to obtain the envelope of the signal. Both separation methods yielded the same results. The same results could also be obtained using conventional electrophysiological techniques (session G98; data not shown, see also Supplementary Information). Specifically, intracortical recordings were carried out with a 4 × 4 array of microfibre electrodes (quartz Pt₉₀W₁₀; 80 μm shaft diameter, spacing of 250 μm centre to centre, impedance 250–750 kΩ, at 500 Hz). The signal was amplified and filtered using BAK amplifiers (MMRS-1S).

Elimination of residual interference by signal processing

The techniques of interference reduction (see above) ensured a non-saturated, measurable signal, but one that was still contaminated with a certain amount of gradient interference. The elimination of the residual interference was accomplished by applying a standard mathematical, dimensionality reduction method—the principal component analysis (PCA) technique. The data were initially realigned to the slice-selection pulse (signifying the beginning of collection of an image of a single K-space segment for single- or multi-shot acquisitions) and subsequently reshaped into an *N* by *M* matrix, where *N* was the number of segments and *M* was the number of data points acquired while digitizing the physiology signal. Analysis with a personal computer of such data and elimination of those principal components that best correlated with the directly recorded interference (obtained from the current-monitor output of the gradient amplifiers) resulted in a ‘clean’ signal as shown in Fig. 7. The quality of the de-noising procedure was tested with artificial signals, and also by comparing the spectra of de-noised signals with signals acquired in the same session (interleaved) without the gradient alternations. We found no difference in their spectra.

Received 21 May; accepted 19 June 2001.

- Ogawa, S. & Lee, T. M. Magnetic resonance imaging of blood vessels at high fields: *in vivo* and *in vitro* measurements and image simulation. *Magn. Reson. Med.* **16**, 9–18 (1990).
- Bandettini, P. A., Wong, E. C., Hinks, R. S., Tikofsky, R. S. & Hyde, J. S. Time course EPI of human brain function during task activation. *Magn. Reson. Med.* **25**, 390–397 (1992).
- Frahm, J., Bruhn, H., Merboldt, K. D. & Hancic, W. Dynamic MR imaging of human brain oxygenation during rest and photic stimulation. *J. Magn. Reson. Imaging* **2**, 501–505 (1992).
- Menon, R. S. *et al.* Functional brain mapping using magnetic resonance imaging. Signal changes accompanying visual stimulation. *Invest. Radiol.* **27**, (Suppl.) 53 (1992).
- Kwong, K. K. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl Acad. Sci. USA* **89**, 5675–5679 (1992).
- Menon, V., Ford, J. M., Lim, K. O., Glover, G. H. & Pfefferbaum, A. Combined event-related fMRI and EEG evidence for temporal-parietal cortex activation during target detection. *NeuroReport* **8**, 3029–3037 (1997).
- Krakow, K. *et al.* EEG recording during fMRI experiments: image quality. *Hum. Brain Mapp.* **10**, 10–15 (2000).

8. Krakow, K. *et al.* EEG-triggered functional MRI of interictal epileptiform activity in patients with partial seizures. *Brain* **122**, 1679–1688 (1999).
9. Bonmassar, G., Anami, K., Ives, J. & Belliveau, J. W. Visual evoked potential (VEP) measured by simultaneous 64-channel EEG and 3T fMRI. *NeuroReport* **10**, 1893–1897 (1999).
10. Hess, A., Stiller, D., Kaulisch, T., Heil, P. & Scheich, H. New insights into the hemodynamic blood oxygenation level-dependent response through combination of functional magnetic resonance imaging and optical recording in gerbil barrel cortex. *J. Neurosci.* **20**, 3328–3338 (2000).
11. Bonhoeffer, T. & Grinvald, A. *Brain Mapping, The Methods* (eds Toga, A. W. & Mazziotta, J. C.) 55–97 (Academic, New York, 1996).
12. Schmitt, F., Stehling, M. K. & Turner, R. *Echo-Planar Imaging: Theory, Technique and Application* (Springer, Berlin, 1998).
13. Legatt, A. D., Arezzo, J. & Vaughan, H. G. J. Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. *J. Neurosci. Methods* **2**, 203–217 (1980).
14. Freeman, W. J. *Mass Action in the Nervous System* (Academic, New York, 1975).
15. Mitzdorf, U. Properties of the evoked potential generators: current source-density analysis of visually evoked potentials in the cat cortex. *Int. J. Neurosci.* **33**, 33–59 (1987).
16. Juergens, E., Guettler, A. & Eckhorn, R. Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical- and EEG-potentials, but not in human EEG. *Exp. Brain Res.* **129**, 247–259 (1999).
17. Eckhorn, R. & Thomas, U. A new method for the insertion of multiple microprobes into neural and muscular tissue, including fiber electrodes, fine wires, needles and microsensors. *J. Neurosci. Methods* **49**, 175–179 (1993).
18. Juergens, E., Eckhorn, R., Frien, A. & Woelbern, T. *Brain and Evolution* 418 (Thieme, Berlin, 1996).
19. Hu, X., Le, T. H. & Ugurbil, K. Evaluation of the early response in fMRI in individual subjects using short stimulus duration. *Magn. Reson. Med.* **37**, 877–884 (1997).
20. Logothetis, N. K., Guggenberger, H., Peled, S. & Pauls, J. Functional imaging of the monkey brain. *Nature Neurosci.* **2**, 555–562 (1999).
21. Maloney, D. & Grinvald, A. Interactions between electrical activity and cortical microcirculation revealed by imaging spectroscopy: implications for functional brain mapping. *Science* **272**, 551–554 (1996).
22. Buxton, R. B., Wong, E. C. & Frank, L. R. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magn. Reson. Med.* **39**, 855–864 (1998).
23. Frahm, J., Kruger, G., Merboldt, K. D. & Kleinschmidt, A. Dynamic uncoupling and recoupling of perfusion and oxidative metabolism during focal brain activation in man. *Magn. Reson. Med.* **35**, 143–148 (1996).
24. Kruger, G., Kleinschmidt, A. & Frahm, J. Dynamic MRI sensitized to cerebral blood oxygenation and flow during sustained activation of human visual cortex. *Magn. Reson. Med.* **35**, 797–800 (1996).
25. Eckhorn, R. *et al.* Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biol. Cybern.* **60**, 121–130 (1988).
26. Murthy, V. N. & Fetz, E. E. Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proc. Natl Acad. Sci. USA* **89**, 5670–5674 (1992).
27. Gray, C. M. & Singer, W. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Natl Acad. Sci. USA* **86**, 1698–1702 (1989).
28. Singer, W. Synchronization of cortical activity and its putative role in information processing and learning. *Annu. Rev. Physiol.* **55**, 349–374 (1993).
29. Tallon-Baudry, C., Bertrand, O., Wienbruch, C., Ross, B. & Pantev, C. Combined EEG and MEG recordings of visual 40 Hz responses to illusory triangles in human. *NeuroReport* **8**, 1103–1107 (1997).
30. Joliot, M., Ribary, U. & Llinas, R. Human oscillatory brain activity near 40 Hz coexists with cognitive temporal binding. *Proc. Natl Acad. Sci. USA* **91**, 11748–11751 (1994).
31. Singer, W. Neuronal synchrony: a versatile code for the definition of relations? *Neuron* **24**, 49–65 (1999).
32. Shadlen, M. N. & Movshon, J. A. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* **24**, 67–77 (1999).
33. Sokoloff, L. in *Basic Neurochemistry* (eds Siegel, G., Agranoff, B., Albers, R. W. & Molinoff, P.) 565–590 (Raven, New York, 1989).
34. Magistretti, P. J., Pellerin, L., Rothman, D. L. & Shulman, R. G. Neuroscience—energy on demand. *Science* **283**, 496–497 (1999).
35. Rothman, D. L. *et al.* *In vivo* nuclear magnetic resonance spectroscopy studies of the relationship between the glutamate-glutamine neurotransmitter cycle and functional neuroenergetics. *Phil. Trans. R. Soc. Lond. B* **354**, 1165–1177 (1999).
36. Shulman, R. G. & Rothman, D. L. Interpreting functional imaging studies in terms of neurotransmitter cycling. *Proc. Natl Acad. Sci. USA* **95**, 11993–11998 (1998).
37. Takahashi, S., Driscoll, B. F., Law, M. J. & Sokoloff, L. Role of sodium and potassium ions in regulation of glucose metabolism in cultured astroglia. *Proc. Natl Acad. Sci. USA* **92**, 4616–4620 (1995).
38. Pellerin, L. & Magistretti, P. J. Glutamate uptake into astrocytes stimulates aerobic glycolysis: a mechanism coupling neuronal activity to glucose utilization. *Proc. Natl Acad. Sci. USA* **91**, 10625–10629 (1994).
39. Sibson, N. R. *et al.* Stoichiometric coupling of brain glucose metabolism and glutamatergic neuronal activity. *Proc. Natl Acad. Sci. USA* **95**, 316–321 (1998).
40. Pan, J. W. *et al.* Spectroscopic imaging of glutamate C4 turnover in human brain. *Magn. Reson. Med.* **44**, 673–679 (2000).
41. Rees, G., Friston, K. & Koch, C. A direct quantitative relationship between the functional properties of human and macaque V5. *Nature Neurosci.* **3**, 716–723 (2000).
42. 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).
43. Sclar, G., Maunsell, J. H. R. & Lennie, P. Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res.* **30**, 1–11 (1990).
44. Boynton, G. M., Demb, J. B., Glover, G. H. & Heeger, D. J. Neuronal basis of contrast discrimination. *Vision Res.* **39**, 257–269 (1999).
45. Boynton, G. M., Engel, S. A., Glover, G. H. & Heeger, D. J. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* **16**, 4207–4221 (1996).
46. Rainer, G., Augath, M., Trinath, T. & Logothetis, N. K. Nonmonotonic noise tuning of BOLD fMRI signal to natural images in the visual cortex of the anesthetized monkey. *Curr. Biol.* **11**, 846–854 (2001).
47. Ugurbil, K. *et al.* Imaging at high magnetic fields: initial experiences at 4 T. *Magn. Reson. Quart.* **9**, 259–277 (1993).
48. Haase, A., Frahm, J., Matthaei, D., Hancic, W. & Merboldt, K.-D. FLASH imaging. Rapid NMR imaging using low flip-angle pulses. *J. Magn. Reson.* **67**, 258–266 (1986).
49. Mansfield, P. Multi-planar image formation using NMR spin echoes. *J. Phys. C* **10** L55–L58 (1977).
50. Gruetter, R. Automatic, localized *in vivo* adjustment of all first- and second-order shim coils. *Magn. Reson. Med.* **29**, 804–811 (1993).

Supplementary information is available on Nature's World-Wide Web site (<http://www.nature.com>) or as paper copy from the editorial office of Nature.

Acknowledgements

We thank D. Leopold, G. Rainer and N. Sigala for reading the manuscript and for many useful suggestions. We also thank H. Mandelkow for writing some of the Matlab code; K. Lamberty for the drawings; D. Blaurock for English corrections and editing; and S. Weber for fine-mechanic work. This research was supported by the Max Planck Society.

Correspondence and requests for materials should be addressed to N.K.L. (e-mail: nikos.logothetis@tuebingen.mpg.de).