Available online at www.sciencedirect.com

SCIENCE DIRECT®

Magnetic Resonance Imaging 24 (2006) 401–409

MAGNETIC RESONANCE IMAGING

The effect of artifacts on dependence measurement in fMRI

Arthur Gretton*, Andrei Belitski, Yusuke Murayama, Bernhard Schölkopf, Nikos Logothetis

Max Planck Institute for Biological Cybernetics, 72076, Tübingen, Germany

Received 2 December 2005; accepted 2 December 2005

Abstract

The study of effective connectivity by means of neuroimaging depends on the measurement of similarity between activity patterns at different locations in the brain, without necessarily presupposing a particular model for this dependence. When these interactions are measured using functional magnetic resonance imaging (fMRI) techniques, however, imaging and physiological artifacts create patterns of dependence that may be unrelated to cortical activity. We demonstrate some of these effects through the measurement of short-range dependencies present in fMRI scans of the primary visual cortex (V1) in the anaesthetized macaque monkey. High-field (4.7 T) fMRI scans were conducted to measure responses based on the blood oxygen level-dependent contrast mechanism, during periods of no sensory stimulation and of visual stimulation with rotating polar-transformed checkerboard gratings. Dependence between the haemodynamic activity at different spatial locations (i.e., different voxels) was measured using correlation, mutual information and functional covariance. Particular attention was paid to understanding the sources of spurious dependence that may be observed during such investigations. Two main effects were detected: (a) short-range correlations introduced by the process of image reconstruction and (b) perturbations in the haemodynamic response caused by breathing. The image reconstruction artifacts were shown to create an artificially high short-range dependence in the readout direction of the scan, and the breathing artifacts caused enhanced short-range dependence in both the readout and phase-encode directions. Additional dependence in the phase-encode direction due to image-ghosting is also possible but will not be discussed in this report, as it can be alleviated by fine adjustment of preemphasis (elimination of eddy currents). A technique is described for removing breathing artifacts, and the effect of breathing on the apparent dependence between voxels is illustrated. The correlation of haemodynamic activity with the stimulus was found to be affected by breathing, although this effect can be neutralised by averaging the haemodynamic responses over many repetitions of the stimulus. Nonetheless, patterns of dependent activity between voxels may be lost in this averaging process, which makes the removal of breathing artifacts necessary if statistical dependence and the study of effective connectivity is the primary aim of an investigation.

Keywords: Dependence; Monkey fMRI; Spatial resolution; Haemodynamic response

1. Introduction

Over the last 100 years, one of the most established concepts in the neurosciences has been the functional specialisation of different brain regions. Often, the pejorative term *cerebral cartography* or even the rather cynical jargon of *novel phrenology* is used to indicate the widespread concept of assigning specific functions to particular cortical areas. The introduction of noninvasive functional magnetic resonance imaging (fMRI) in humans has extended the trend of generating cortical maps into a systematic search for homologies between the human brain areas and those already defined in the much-studied nonhuman primates. Although there is still lively discussion regarding the details of different types of cortical parcellation, there is good agreement that sensory information is organised at different scales and spans different brain areas that work in parallel and hierarchically, and these areas have very dense interconnections. Unravelling the neural events, underlying behaviour is contingent upon understanding both the local-scale neural activity patterns and the global-scale interactions between cortical and subcortical (e.g., thalamus and basal ganglia) structures.

Vision is an extensively studied modality in both humans and nonhuman primates, where both neuroimaging and physiology have been applied. Indeed, more than two dozen different areas have been identified in the visual system of the monkey, on the basis of criteria such as (a) architecture, (b) connectivity and (c) visual topography and/or functional

* Corresponding author. Tel.: +49 7071 601562.
E-mail address: arthur@tuebingen.mpg.de (A. Gretton).

0730-725X/$ – see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.mri.2005.12.036

characteristics (for a review, see Ref. [1]). The number of connections identified between these areas is larger than three hundred, and most of them are reciprocal. Recently, attempts to obtain a metrical representation of this bewildering connectivity in monkeys and cats, by means of nonmetric multidimensional scaling, showed that the cortices of higher mammals are organised into densely intraconnected clusters that reflect their functional specialisation, suggesting that structure and function are indeed closely linked at this gross systems level [2].

In light of such findings, understanding how the activity that spreads over many of these areas yields visual perception, cognition and visually guided action has become a central quest in systems neuroscience. Of particular importance is the development of methods to study this extensive connectivity in longitudinal in vivo experiments that do not require fixed brain, typically precluding histological approaches. For one, processes including learning, plasticity and reorganisation can only be studied in longitudinal experiments where the link between function and structure is investigated and described in detail over long periods of time in the same subject. Moreover, knowledge of in vivo connectivity can provide important subject-specific information for multisite, multielectrode intracortical recordings in combined behavioural and physiology experiments. The value of such information cannot be overstated. If we are to understand how sensory or perceptual representations lead to the awareness of objects or the guidance of action, local information about single-neuron or small neural network activity must indeed be acquired simultaneously from multiple sites that are effectively connected.[1] The massive, reciprocal interconnectivity suggests that only a small fraction of the synapses in large-scale networks are likely to be active at any given behavioural condition. In addition, effective connectivity is likely to change as a result of training, learning or habituation.

A number of invasive or noninvasive approaches have been recently applied to study connectivity in vivo, including the usage of paramagnetic tracers (see, e.g., Ref. [3] and the article by Murayama et al. in this issue), and the development of different theoretical and neuroimaging approaches investigating the organisational principles of the brain, in terms of distributed and coupled interactions among specialised brain system [4–6]. While many fMRI experiments rely on specific models (often generalised linear models [7]) to determine the response to a particular stimulus, this approach is not necessarily optimal if our aim is to characterise the dynamic interactions arising in response to different types of stimulation. In the latter case, a general model-free measure of dependence between the haemodynamic activity measured at different voxels may be preferable. Yet, it then becomes necessary to take account of noise sources in the haemodynamic response, which are concomitant with the signal arising from neural processing. In particular, such noise sources can exhibit strong correlations between different locations, which can result in misleading and irrelevant dependencies being found where no underlying neural basis exists. For example, the low spatiotemporal resolution of most human studies, commonly used image reconstruction methods, or an insufficient pre-emphasis adjustment to suppress ghosting artifacts, can all induce artificial dependence patterns. While the description and understanding of artifacts in fMRI imaging is well established, their removal remains a topic of current research (see, e.g., [8,9]).

Recently, blood oxygen level-dependent (BOLD) fMRI was combined in the nonhuman primate with simultaneous intracortical recording using novel electrophysiological techniques [10,11]. In an attempt to understand the spatial dependence of different intracortical sites and the effect of this dependence in the point-spread image of the BOLD signal for different nominal spatial resolutions, we set out to conduct separate and simultaneous multielectrode and fMRI experiments in high-field (4.7 T) scanners. In the current study, we investigate two sources of spurious dependence observed in these fMRI scans. The first source is an image processing artifact which resembles a Gibbs effect (a short range ripple that decreases with distance) and which increases the apparent dependence between closely spaced voxels in the readout (frequency encoding) direction. We confirm the range and shape of the rippling effect using a scan conducted on a phantom (a cylinder filled with water). The second source is the oscillation in the haemodynamic response at each voxel due to breathing, the strength of which depends on the proximity of the voxels to larger blood vessels. We observe both rippling and breathing effects in fMRI scans of the primary visual cortex (V1) of an anaesthetized macaque monkey, during the presentation of flashing polar gratings: in particular, the breathing artifacts were predictable under our experimental setting, and the effect of their removal on the dependence structure could be measured.

We make use of three different measures of dependence: the correlation, which addresses only second-order interactions, and two more general dependence measures: the mutual information [12] and the constrained functional covariance [13]. We observe that the dependence between voxels as a function of distance between them is substantially affected by both breathing and ripple artifacts, and that an accurate picture of the length scale of interactions in V1 is only obtained when these artifacts are accounted for.

We begin our presentation in Section 2 with a description of the experimental techniques used to acquire our measurements of the BOLD response, followed in Section 3 by an overview of the dependence measures used. In Section 4, we demonstrate the effect of artifacts on the dependence

[1] In general, functional or effective connectivity describes the interdependence in measured activity of brain regions that are either directly or indirectly anatomically interconnected.

measured between voxels as a function of average distance between them, both for a phantom (for which no dependence is to be expected, absent that due to imaging noise), and in the macaque visual cortex. In addition, we show that the removal of breathing artifacts can, in some circumstances, increase the number of voxels that exhibit significant correlation with the stimulus.

2. Methods

A single rhesus monkey was used for the experiments reported in this study. All experiments were approved by the local authorities (Regierungspraesidium) and were in full compliance with the guidelines of the European community (EUVD 86/609/EEC) for the care and use of the laboratory animals.

2.1. Surgery

Details on surgery and anaesthesia are given in previous publications [11,14]. Briefly, after premedication with glycopyrolate (0.01 mg/kg im) and ketamine (15 mg/kg im), a 20-gauge intravenous catheter was introduced into the saphenous vein, and the monitors [HP OmniCare/CMS; electrocardiogram (ECG), noninvasive blood pressure (NIBP), CO_2, SpO_2, temperature] were connected. The monkeys were preoxygenated, and anaesthesia was induced with fentanyl (3 μg/kg), thiopental (5 mg/kg) and succinylcholine chloride (3 mg/kg). Following the intubation of the trachea, the lungs were ventilated using the Siemens Ventilator SV 900 C (Siemens-Elema AB, Sweden), maintaining an end-tidal CO_2 of 33 mmHg and oxygen saturation over 95%. Balanced anaesthesia was maintained with end-tidal 0.35% (0.23 MAC for macaques) isoflurane in air and fentanyl (3 μg kg^{-1} h^{-1}). Muscle relaxation was achieved with mivacurium (5 mg kg^{-1} h^{-1}). Body temperature was kept constant, and lactated Ringer's solution was given at a rate of 10 ml kg^{-1} h^{-1}. Intravascular volume was maintained by administering colloids (hydroxyethyl starch, 30–50 ml over 1–2 min, as needed). Emergence from anaesthesia was typically without complications and lasted an average of 30 min. The paralytic and fentanyl were stopped, and ventilation was lowered to stimulate spontaneous breathing. When spontaneous respiration was assured and the CO_2 was below 40 mmHg, the trachea was extubated. Temperature, ECG, NIBP, CO_2 and SpO_2 were monitored throughout surgery.

2.2. MRI data collection and analysis

Measurements were made on a vertical 4.7-T scanner with a 40-cm diameter bore (Biospec 47/40v, Bruker Medical, Ettlingen, Germany). The system was equipped with a 50-mT/m (180-μs rise time) actively shielded gradient coil (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. Customised small radiofrequency (RF) coils of diameter 30–80 mm were used, which were optimised for increased sensitivity over a given region of interest (ROI) such as a portion of the primary visual cortex or regions of areas V2, V3, V4 and V5 (the mediotemporal visual area). The results in this paper pertain solely to the activation of the striate cortex, V1. 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. Gradient echo fast imaging with a repetition time (TR) of 1800 ms, echo time (TE) of 10.4 ms and flip angle (FA) of 30° was used to obtain T1-weighted anatomical slices (0.5-mm thickness) of a 512×256 matrix (0.1875×0.1875-mm resolution).

To study the time course of the BOLD signal, high spatial resolution functional signals were acquired as a series of 128×64×2 (0.75×0.75×2-mm resolution) volumes by using blipped phase encoding echo planer imaging (EPI) with a TR of 125 ms, TE of 20 ms, FA of 40°, and EPI-zero-phase of 8.11 ms. In order to minimise the effects of inflow and of large drainage vessels, we consistently used FAs 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 was typically performed with a hard pulse of 50-μs duration and a receiver bandwidth of 50 kHz. Depending on the RF coil, either global or local volume shimming (FASTMAP [15]) was used. The line width (full width at half height) at the end of the shimming procedure ranged from 30–70 Hz for the entire scanned volume. Together with the image acquisition, both the plethysmogramm (detection of volume changes due to arterial pulsations, roughly corresponding to an ECG signal) and the flow signal of the ventilator were digitised at 250 Hz and stored for later evaluation and removal of physiological artifacts from the BOLD time courses.

Each scan was normalised (ratio normalisation) by dividing each pixel's intensity value by the mean intensity of a tracer-free region in the scan and multiplying the result by 5000. This procedure is immune to the considerable differences in the average intensity that can arise during data collection as a result of factors such as differences in receiver gain or in the quality factor, Q, of coils.

3. Dependence measurement

In this section, we introduce three measures of dependence between random variables, which we will apply in the next section to the problem of determining whether the measured haemodynamic activity at different voxels is related. Suppose we have a pair of random variables $X;Y$, with a joint density $p(X;Y)$, and we wish to determine whether there is any dependence between the two. For our purposes, dependence is defined as the density not factorising as a product of the marginals, i.e., $p(X;Y) \neq p(X)p(Y)$. In practice, we are not given the density itself, but a sample of m pairs $Z := \{(x_1; y_1),\ldots,(x_m;y_m)\}$ distributed indepen-

dently and identically[2] according to $p(X;Y)$: any estimate of dependence we make is therefore approximate and becomes more accurate as we obtain more samples.

The simplest measure of dependence we consider is the correlation, defined as

$$\mathrm{corr}(X,Y) = \frac{cov(X,Y)}{\sqrt{var(X)var(Y)}}$$

where $\mathrm{cov}(X;Y) = E(XY) - E(X)E(Y)$ and $\mathrm{var}(X) = E(X^2) - (E(X))^2$. An empirical estimate of the correlation may readily be obtained by replacing the expectations with their empirical estimates. That said, the correlation only accounts for second-order effects: two signals may be dependent and yet have zero correlation. A simple example is the correlation between a sine and a cosine with equal frequency: the two are clearly dependent, yet their correlation over a sufficiently long time approaches zero.

We next describe the mutual information [12]. Given two random variables $X;Y$ with joint density $p(X;Y)$, the mutual information is

$$I(X,Y) = \int_x \int_y p(X,Y)\log\left(\frac{p(X,Y)}{p(X)p(Y)}\right)dXdY$$

where $p(X)$ and $p(Y)$ are the marginal densities corresponding to the joint density $p(X;Y)$. This quantity can be interpreted as the information that the random variables X and Y share; or in other words, the degree to which one can predict Y from X (and vice versa). The mutual information can be expressed in units of bits or nats, depending on the base of the logarithm. A useful property of the mutual information is that it is zero if and only if the random variables X and Y are independent, in which respect, it is more general than the correlation. That said, the empirical estimation of the mutual information is considerably more challenging than the correlation — in the present study, we use the approach of Ref. [17].

Finally, we define the constrained functional covariance (COCO), which (like the mutual information) is a general measure of dependence. This approach to dependence measurement was proposed by Ref. [18], and statistical tests of dependence based on a related quantity (the functional correlation) were obtained in Ref. [19]. Here, we follow the presentation in Ref. [13]. Given function classes $F;G$ and the density $p(X;Y)$, we define COCO as

$$\mathrm{COCO}(p(X,Y);F,G) := \sup_{f\in F, g\in G}[cov(f(X),g(Y))]. \quad (3.1)$$

When $\mathcal{F}$ and $\mathcal{G}$ are universal reproducing kernel Hilbert spaces,[3] with F and G their respective unit balls, then $\mathrm{COCO}(P_{x;y};F;G)$ is guaranteed to be zero if and only if X and Y are independent. The empirical estimate of COCO is straightforward [13], and is given by

$$\mathrm{COCO}(Z;F;G) = \frac{1}{m}\sqrt{||\tilde{K}\tilde{L}||_2},$$

where the norm $\|\cdot\|_2$ denotes the largest singular value, and

$$\tilde{K} := HKH,\ \tilde{L} := HLH,$$

where

$$H = I - \frac{1}{m}1_m 1_m^T, \quad (3.2)$$

1_m is an $m \times 1$ vector of ones, K is a matrix with $(i;j)$th entry

$$k(x_i,x_j) = \exp\left(-\frac{1}{2}||x_i - x_j||^2\right)$$

and L has $(i;j)$th entry $k(y_i;y_j)$.

4. Dependence artifacts observed in fMRI scans

Having introduced three alternative measures for dependence between random variables, we now describe the impact of structured interference in fMRI scans on estimates of dependence between the haemodynamic signal at different locations (we define noise as being any aspect of the signal that is not related to the haemodynamic response contingent on neural activity, although we clearly do not attempt to account for all possible such noise sources). In particular, we demonstrate that the dependence is greatly affected by structured noise sources, which (unless properly accounted for) can give a severely misleading picture of the extent and size of this dependence.

This section is divided into two parts corresponding to the two sources of structured noise that we investigate. The first noise source is a spatial rippling effect, which arises in the process of image reconstruction and resembles a Gibbs artifact (i.e., a decaying oscillation that might arise from the truncation of high spatial frequency terms: see Ref. [21]). The second effect is a sinusoidal oscillation in the haemodynamic response due to breathing. In our first experiment, described in Section 4.1, we investigated the rippling effect in isolation by scanning a phantom and observing the short-range dependence patterns thus revealed (these can only be due to scanning artifacts, since there is no biological signal present). In Section 4.2, we describe our second experiment, in which we measured the dependence between voxels in the V1 region of the macaque visual cortex, arising from a flashing polar grating. In particular, we show how this dependence is affected by both breathing

[2] Note that this assumption does not correspond to the time series observed at each voxel in an fMRI scan, which are clearly random processes, and not iid samples. There exist more sophisticated measures of dependence that take this into account, such as the information rate [12,16], but these are beyond the scope of the present work.

[3] See Ref. [20] for a definition of universality: the mathematical details are beyond the scope of the present discussion.

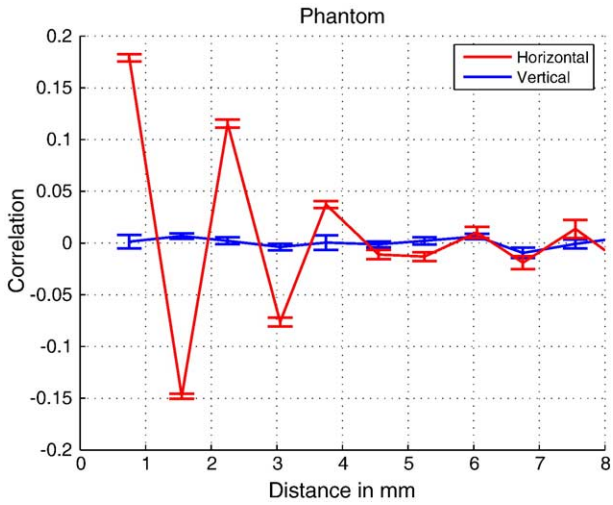

Fig. 1. Demonstration of correlation vs. distance in a phantom recording. We observe that the vertical (phase encoding) dependence vs. distance structure is effectively flat, whereas the horizontal (readout) structure shows strong correlations (that decay with distance). This is despite the fact that the time series recorded at each voxel show no obvious temporal or spectral structure on visual inspection.

artifacts and the spatial rippling. Moreover, we find that the correlation of the haemodynamic response with the stimulus can be decreased by the presence of breathing artifacts.

4.1. Artifacts observed in scans of a phantom

To investigate artifacts arising solely from the data processing techniques used in our fMRI recordings, we performed scans on a phantom (a cylinder filled with water), using the same settings as were applied in measuring the BOLD response (as described in the next section). For this experiment, we measured dependence using only the correlation, since this gives the clearest picture of the rippling behaviour. The variation in correlation between voxels was studied as a function of average distance between voxels (in other words, we grouped together all pairs of voxels with an equal distance from each other; we then clustered these pairs so as to draw together voxel pairs with similar distances). We see in Fig. 1 that there is a strong correlation structure in the horizontal direction (the readout or frequency-encoded direction), which resembles a Gibbs artifact (we do not seek in the present study to explain the source of this behaviour but only to note its effect on the apparent dependence between voxels. That said, one possible source of this artifact is an incomplete or inadequate ramp-sampling correction, which can cause correlation patterns akin to those observed). The vertical (phase encoded) direction is unaffected by this particular noise source. We expect, therefore, that an accurate picture of dependence vs. distance will not be obtained in the readout direction if the distances are small enough for the dependence to be partly due to this rippling.

4.2. Dependence structure in the macaque visual cortex

Our second experiment consisted in measuring haemodynamic activity in the primary visual cortex (V1) of the anaesthetised macaque. Since our intention is to obtain model-free estimates of dependence between the neural activity underlying the BOLD response, we follow Ref. [13] in applying the three methods described in Section 3 to find the degree of similarity in the signal at pairs of voxels. There are two main differences with respect to our previous work in Ref. [13], however. First, we treat the horizontal and vertical dependencies separately, which allows us to isolate the effect of the rippling artifact observed in our scan of the phantom. Second, we used a ashing polar grating as stimulus (rather than a film clip), which resulted in a simple periodic activation across much of V1: this should, in turn, cause a straightforward dependence structure to be observed between the voxels (and a prompt decay in this dependence as we move out of those parts of V1 correlated with this stimulus). Specifically, the monkey was exposed to a

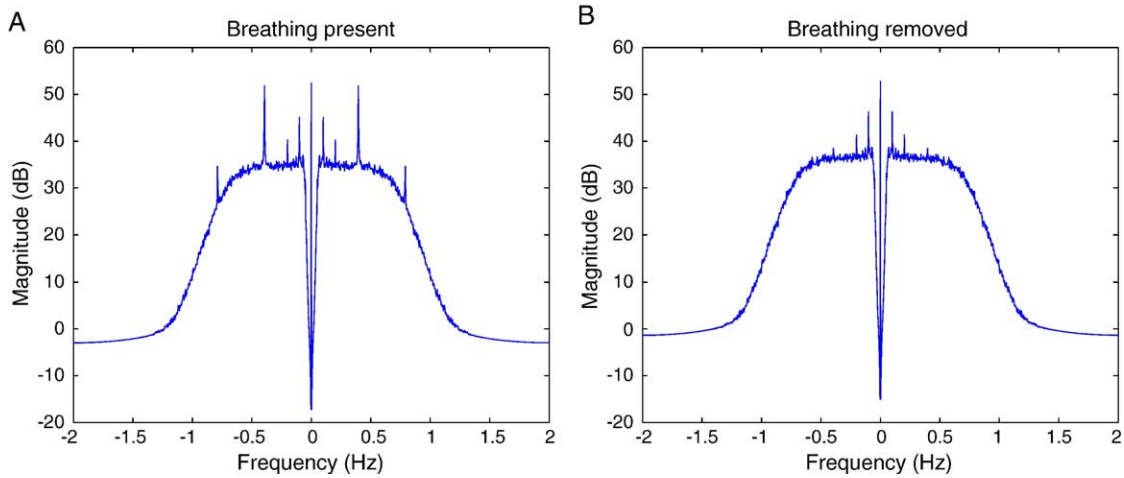

Fig. 2. Average spectrum across voxels in V1 when the polar ash stimulus is used (at 0.1 Hz). A response at the stimulus frequency and the first harmonic is clearly present, as are additional peaks due to breathing (in the left hand plot).

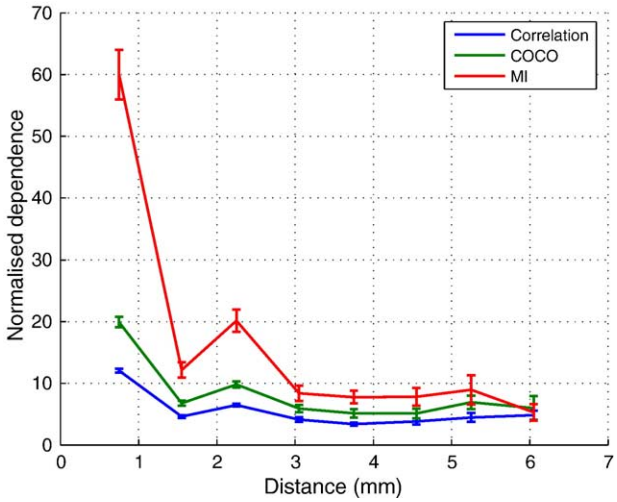

Fig. 3. Dependence vs. distance for horizontal distances only. Breathing artifacts were removed, but the rippling effect contributes significantly to the overall dependence at small distances.

periodic cycling of 2 s without stimulus, 3 s of stimulus, and 5 s of no stimulus. The ROI in our scan representing V1 contained a total of 95 voxels, and the total signal duration was 320 s (1280 samples at 4 Hz). The results presented are an aggregate over five such experiments.

The observed fMRI signals were contaminated with a breathing component. Since the macaque monkey was under general anaesthetic during data acquisition, breathing was mechanically assisted and had a constant frequency of approximately 0.4 Hz. We modelled this breathing as being of constant amplitude and linearly superposed on the haemodynamic response. This model is motivated by the narrowness of the spectral peaks at the breathing frequency and harmonics, which suggests that any amplitude modulation of the breathing signal is of very low frequency and can be assumed effectively constant. Thus, while we could not directly recover the true breathing contamination at each voxel, we were able to use the decrease in the spectral peak at the breathing frequency, averaged across all voxels, as a measure of success in removing the breathing artifacts. Harmonics at integer multiples of the fundamental frequency were modelled in the same way. The exact frequency of the breathing signal was found by averaging the spectrum over all voxels, and the phase at each voxel was chosen to maximise the projection in the time domain of the breathing sinusoid. Only voxels near large blood vessels were contaminated by the breathing signal and, thus, a threshold test was applied to the spectrum at each voxel, to test whether a substantial breathing component was visible. Where breathing was present, a sinusoid of corresponding frequency and phase was projected out in the time domain (thus also removing the associated frequency domain sidelobes caused by nite signal duration). The same procedure was used to remove the first two harmonics. We did not band-pass filter the signal to remove the breathing, as this would have eliminated a greater portion of the spectral components due to the haemodynamic activity. The average spectrum across all recorded V1 voxels is plotted in Fig. 2, both before and after the breathing removal. The spectral peaks due to the breathing are clearly present, as are lower frequency peaks attributable to the stimulus.

We constructed dependence vs. distance plots in the same manner as for the phantom. As dependence measures between pairs of voxels, we used the absolute value of the cross correlation between voxels, the mutual information and COCO. In addition, we subtracted a baseline dependence from each of the dependence measures, which was obtained by averaging the dependence between the V1 region and a test region of the brain, the latter being unrelated to visual processing. The dependence amplitudes were then divided by the standard deviation in the average dependence between V1, and this test region to obtain a normalised dependence. Dependence as a function of horizontal distance (where scanning artifacts dominate the dependence structure) is plotted in Fig. 3. Dependence as a

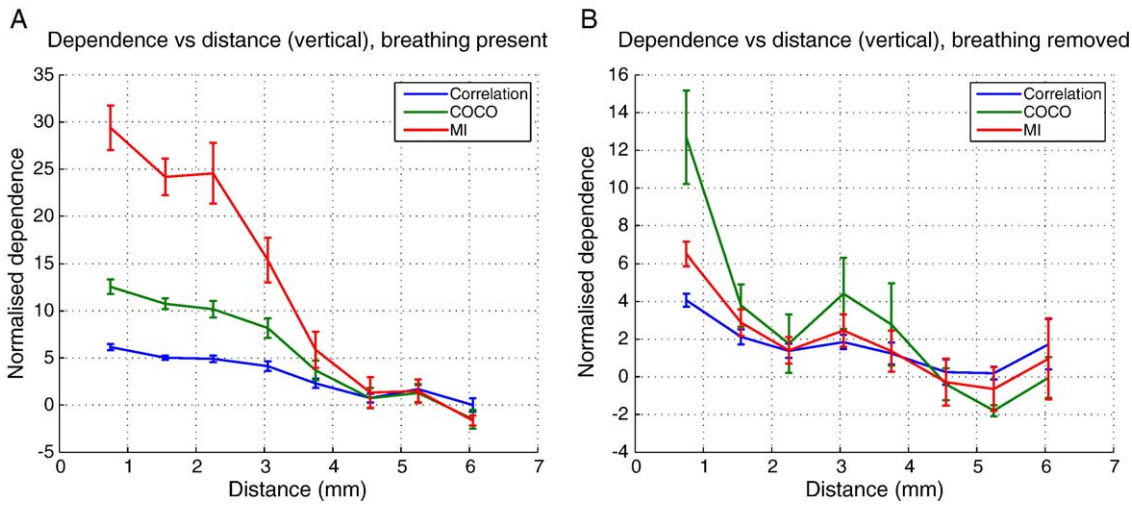

Fig. 4. Dependence vs. distance for vertical distances only. Plots are given both before (left) and after (right) breathing artifact removal.

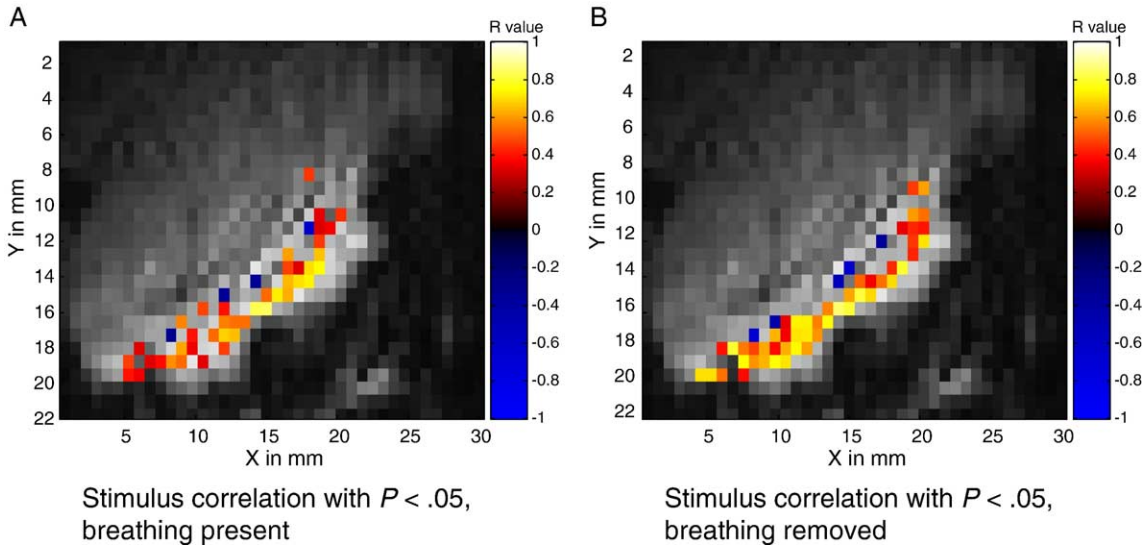

Fig. 5. Correlation of voxels with the stimulus (the correlation was obtained following a convolution of the stimulus waveform with the expected haemodynamic response). In this case, the haemodynamic response was averaged over five repetitions of the stimulus. The left hand plot shows the correlations with the stimulus, in the absence of breathing removal, and using the threshold $P<.05$. The right hand plot illustrates the correlations, with the same threshold, following breathing removal.

function of vertical distance, both before and after breathing artifact removal, is plotted in Fig. 4.

Comparing the dependence measures before and after breathing removal shows significant effects of respiratory artifacts on the dependence vs. distance curves: in particular, breathing effects boost the apparent dependence at short distances and caused the dependence to decay more slowly with increasing distance (the breathing artifacts being a source of additional similarity). This finding suggests extreme caution for studies in humans, in which respiration-induced signals cannot easily be modelled due to low temporal sampling rates, as well as variable respiration frequency and amplitude.

It is also instructive to compare the horizontal and vertical dependence structures. In particular, we note that in the horizontal direction, the dependence is greater and over a longer range than we observe in the vertical direction. While the orientation of V1 means that we obtain longer contiguous regions of voxels correlated with the input (and, hence, large dependence over a greater distance), the fact

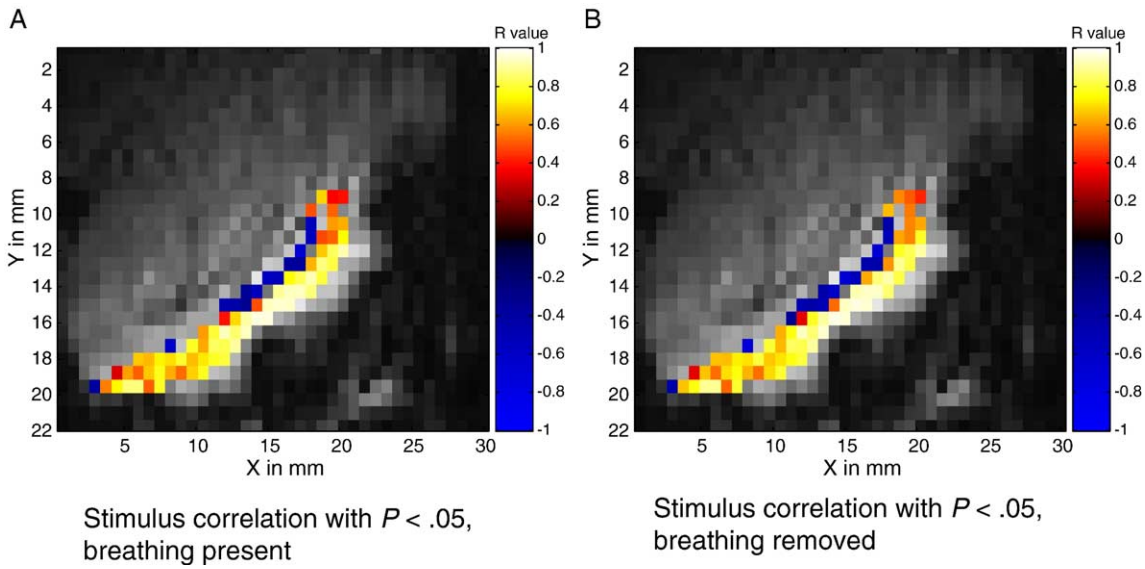

Fig. 6. Correlation of voxels with the stimulus (the correlation was obtained following a convolution of the stimulus waveform with the expected haemodynamic response). In this case, the haemodynamic response was averaged over 150 repetitions of the stimulus. The left hand plot displays the correlations in the absence of breathing removal and using the threshold $P<.05$. The right hand plot shows the correlations, with the same threshold, following breathing removal.

remains that a large proportion of the dependence observed, particularly at small intervoxel distances, is due to scanning artifacts of the sort we obtained with the phantom.

Finally, it is of interest to investigate how correlations of the voxels with the stimulus are affected by breathing artifacts. A plot of these correlations is given in Fig. 5, where the correlations with the stimulus signal (convolved with a model haemodynamic response) are given both before and after breathing is removed. The correlation was not computed with respect to the raw haemodynamic response: rather, the time series at each voxel was averaged over five repetitions of the stimulus, so as to (somewhat) reduce noise. We observe that removal of the breathing causes us to recover a greater number of voxels significantly correlated with the stimulus (at the significance level $P = .05$), and that the correlation generally increases when breathing artifacts are eliminated. We note, however, that averaging the haemodynamic response over a greater number of repetitions of the stimulus also reduces the effect of breathing on the correlations. We illustrate this effect in Fig. 6, where an average response at each voxel is taken over 150 presentations of the polar grid. Indeed, the breathing occurs at higher frequencies than the stimulus, and with a different phase during each repetition: when taken over sufficient repetitions, the average of these sinusoidal signals vanishes. That said, this averaging process can be misleading when attempting to recover intervoxel dependence, which might be unique to each trial. Thus, it is preferable to remove the breathing signal, rather than averaging over it, in cases where the data acquisition rate and uniformity of the breathing over time (such as during anaesthesia) make this possible.

5. Conclusion

The global context of this study is the statistical analysis of dependence in neural activity in the nonhuman primate, using both fMRI and electrophysiology (following MRI-driven electrode implantation), to determine effective connectivity. Its specific aim has been to identify and eliminate artifacts in the observed BOLD response, which can generate dependence patterns unrelated to neural connectivity or to stimulus-induced covariation of activity in different cortical regions. In particular, both respiration and short range Gibbs-type oscillations (the latter being observed in the readout direction and arising as a result of the image reconstruction procedure) can yield a misleading picture of the size and extent of these dependence patterns. While the effect of the Gibbs artifact on the dependence can be circumvented by measuring dependence only between voxels separated in the phase-encoded direction, this places an artificial limit on the configuration of the anatomical structures, for which dependence may be investigated. It is therefore of interest to determine the source of this rippling artifact in the imaging process, with a view to its removal or to compensation of the effect on dependence that it induces.

In addition, we have shown that the removal of breathing artifacts can increase the correlation of voxels with the stimulus and that averaging the haemodynamic response over many repetitions of the stimulus also counters the effect of the (relatively) high frequency breathing signal. That said, it is not practical to average over a large number of repetitions in experiments where dynamic natural stimuli (such as film clips) are used. In these cases, and if the breathing can be maintained at a constant frequency and amplitude (as ensured here due to the anaesthetized subject), breathing removal can be accomplished using the projection approach outlined above and in Ref. [13]. We note that cardiac activity presents an additional source of structured noise in the haemodynamic response. This is much more difficult to treat than the breathing due to its complex shape (it is far from sinusoidal, and the spectral peak is correspondingly broad), and due to it being aliased (the heartbeat occurs at a higher frequency than the Nyquist rate attainable in most fMRI scans; the waveform resulting from the aliased cardiac signal is about 2.5–3 Hz, while the sampling rate is 4 Hz). Nonetheless, the removal of this signal may be feasible, given that we have a separate high-frequency recording of the heartbeat via the plethysmogram.

Acknowledgment

We would like to thank Jozien Goense, Mark Augath and Michael Deelwater for helpful discussions. This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence (IST-2002-506778) and by the Max Planck Society.

References

[1] Felleman D, Van Essen D. Distributed hierarchical processing in the primate cerebral cortex. Cereb Cortex 1991;1:1–47.
[2] Hilgetag C, Burns G, O'Neill M, Scannell J, Young M. Anatomical connectivity defines the organization of clusters of cortical areas in the macaque monkey and the cat. Philos Trans R Soc Lond B Biol Sci 2000;355:91–110.
[3] Saleem K, Pauls J, Augath M, Trinath M, Prause B, Logothetis N. Magnetic resonance imaging of neuronal connections in the macaque monkey. Neuron 2002;34:685–700.
[4] Mechelli A, Price C, Noppeney U, Friston K. A dynamic causal modeling study on category effects: bottom-up or top-down mediation? J Cogn Neurosci 2001;15:925–34.
[5] Mechelli A, Penny W, Price C, Gitelman D, Friston K. Effective connectivity and intersubject variability: using a multisubject network to test differences and commonalities. Neuroimage 2002;17:1459–69.
[6] Friston K, Harrison L, Penny W. Dynamic causal modelling. Neuroimage 2003;19(4):1273–302.
[7] McCullagh P, Nelder J. Generalized linear models. 2nd ed. Boca Raton (FL): CRC Press; 1989.
[8] Bernstein M, King K, Zhou X. Handbook of MRI pulse sequences. New York: Elsevier; 2004.
[9] Haacke E, Brown R, Thomson M, Venkatesan R. Magnetic resonance imaging physical principles and sequence design. New York: Wiley; 1999.

[10] Logothetis N, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. Nature 2001;412:150–7.
[11] Logothetis N, Guggenberger H, Pauls J, Peled S. Functional imaging of the monkey brain. Nat Neurosci 1999;2:555–62.
[12] Cover TM, Thomas JA. Elements of information theory. New York: John Wiley and Sons; 1991.
[13] Gretton A, Smola A, Bousquet O, Herbrich R, Belitski A, Augath M, et al. Kernel constrained covariance for dependence measurement. AISTATS, vol. 10, 2005. URL for proceedings: http://www.gatsby.ucl.ac.uk/aistats/Alabst.htm.
[14] Logothetis N, Merkle H, Augath M, Trinath T, Ugurbil K. Ultra-high resolution fMRI in monkeys with implanted RF coils. Neuron 2002; 35(2):227–42.
[15] Gruetter R. Automatic, localized in vivo adjustment of all first- and second-order shim coils. Magn Reson Med 1993;29:804–11.
[16] Strong S, Koberle R, de Ruyter van Steveninck R, Bialek W. Entropy and information in neural spike trains. Phys Rev Lett 1998;80(1): 197–200.
[17] Darbellay GA, Vajda I. Estimation of the information by an adaptive partitioning of the observation space. IEEE Trans Inf Theory 1999; 45(4):1315–21.
[18] Rényi A. On measures of dependence. Acta Math Acad Sci Hung 1959;10:441–51.
[19] Dauxois J, Nkiet GM. Nonlinear canonical analysis and independence tests. Ann Stat 1998;26(4):1254–78.
[20] Steinwart I. On the influence of the kernel on the consistency of support vector machines. JMLR 2002;2:67–93.
[21] Oppenheim A, Willsky A, Nawab S. Signals and systems. 2nd ed. New York: Prentice Hall; 1996.