

The chronoarchitecture of the human brain—natural viewing conditions reveal a time-based anatomy of the brain

Andreas Bartels* and Semir Zeki

Wellcome Department of Imaging Neuroscience, University College London, London, UK

Received 21 October 2003; revised 5 January 2004; accepted 5 January 2004

A dominant tendency in cerebral studies has been the attempt to locate architecturally distinct parts of the cortex and assign special functions to each, through histological, clinical or hypothesis-based imaging experiments. Here we show that the cerebral cortex can also be subdivided into different components temporally, without any a priori hypotheses, based on the principle of functional independence. This states that distinct functional subdivisions have activity time courses (ATCs) that are, if not independent, at least characteristic to each when the brain is exposed to natural conditions. To approach a time-based anatomy experimentally, we recorded whole-brain activity using functional magnetic resonance imaging (fMRI) and analyzed the data with independent component analysis (ICA). Our results show that a multitude of cortical areas can be identified based purely on their characteristic ATCs during natural conditions. We demonstrate that a more ‘rich’ stimulation (free viewing of a movie) leads to more areas being activated in a specific way than conventional stimuli, allowing for a more detailed dissection of the cortex into its subdivisions. We show that stimulus-driven functionally specialized areas can be identified by intersubject correlation even if their function is unknown. Chronoarchitectonic mapping thus opens the prospect of identifying previously unknown cortical subdivisions based on natural viewing conditions by exploiting the characteristic temporal ‘fingerprint’ that is unique to each.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Chronoarchitecture; fMRI; Time-based anatomy; Independent component analysis; ICA; Movie; Visual cortex; Auditory cortex; Connectivity

Introduction

A major thrust of cerebral studies has been the attempt to subdivide the cerebral cortex into anatomically distinct compartments and assign specific functions to each. Since function is related to structure, and is indeed defined by it (Vogt and Vogt, 1919), anatomical and functional methods should ultimately be

equal to the task of delineating subdivisions in the cortex. Cytoarchitectonic maps (Brodmann, 1909) have provided crude cerebral subdivisions where single areas, such as area 18, have subsequently been shown to contain several distinct, functionally specialized, areas (Zeki, 1978). Cytochrome oxidase maps (Livingstone and Hubel, 1984) have revealed more local detail, but only in a few areas, for example the subdivisions within V2 (DeYoe and Van Essen, 1985; Hubel and Livingstone, 1987; Shipp and Zeki, 1985). Myeloarchitectonic maps (Flechsig, 1901) have a more direct relation to function, as they imply, with hindsight, differences in signal arrival times and internal processing speed. Even within the modality of vision, these can vary widely, which is reflected in differential signal arrival times and perceptual delays (Beckers and Zeki, 1995; Moutoussis and Zeki, 1997; Schmolesky et al., 1998; Zeki and Bartels, 1998). Just as histologically derived maps depend on the markers used, maps derived from conventional fMRI experiments depend on the stimuli and prior hypotheses (Fracowiak et al., 1997; Friston et al., 1995b). Here we address the question of whether one can generate maps of the cerebral cortex without any prior knowledge, neither of its anatomical constitution nor about its functional response. The marker we propose to use is strictly the time course of activation (ATC) of different cortical components during natural conditions, which is, as we will postulate, characteristic to each functional and structural subdivision of the cortex, just like a fingerprint. We argue that this is a logical consequence of the fact that distinct areas are specialized to process distinct features (Zeki, 1978), and that distinct (internal or external) features can vary independently over time in natural conditions. We therefore formulate the hypothesis that each functionally specialized area has an activity time course (ATC) that is, if not independent, at least characteristic to it when the brain is exposed to natural conditions. We refer to above hypothesis as to the principle of functional independence (Bartels and Zeki, 2004b). According to it, the more complex and rich the tasks, the more subdivisions should be activated in a distinct manner. It should therefore be possible, at least theoretically, to dissect the cortex ‘blindly’ into its functional subdivisions, based entirely on temporal markers. To achieve this, we analyzed two data sets. The first was collected while subjects participated in a conventional task, which allowed us to identify specialized areas using conventional methods (Friston et al., 1995b) and compare the results to those obtained using the ‘blind’, data-driven approach of independent components analysis (ICA). The latter has been shown previously

* Corresponding author. Max Planck Institute for Biological Cybernetics, Spemannstrasse 38, Tuebingen 72076, Germany. Fax: +49-7071-601-652.

E-mail address: andreas.bartels@tuebingen.mpg.de (A. Bartels).

Available online on ScienceDirect (www.sciencedirect.com.)

to separate artifacts and regions involved in separate tasks without a priori hypotheses (Bartels and Zeki, 2000a; Calhoun et al., 2001a; Friston, 1998; McKeown, 2000a; McKeown and Sejnowski, 1998; McKeown et al., 1998a; Zeki and Bartels, 1999). Second, we stimulated the brains of eight subjects with the 22 opening minutes of the movie *Tomorrow Never Dies*. We expected the dynamics and complexity of this stimulus to be more likely to stimulate a multitude of functionally specialized areas differentially over time. The results of the first analysis confirmed that ICA can isolate functionally specialized areas, which were also identified using statistical parametric mapping (SPM) (Friston et al., 1995b), and revealed that distinct areas exhibit quite distinct ATCs during epoch stimulation. The analysis of the second study revealed that even during natural viewing, distinct areas exhibit characteristic ATCs. Further, it revealed that the parcellation of the cortex by ICA was consistent across different brains, and that corresponding regions of the brain in distinct subjects had corresponding ATCs, which can be exploited to recognize functional subdivisions across different brains. Furthermore, we show that distinct subdivisions of the brain have activity time courses that vary with a very high degree of independence from each other, which may be explained by the temporal independence of features in the stimulus (Bartels and Zeki, 2003). Finally, a comparison between the epoch and the free viewing data revealed that many more areas are activated with distinct ATCs during natural viewing than during conventional stimulation.

Methods

ICA applied to fMRI: background and motivation

Based on information theory, ICA is the most powerful method for unmixing or decomposing linear mixtures of independent sources, which need not be known a priori (Bell and Sejnowski, 1995). The (mixed) data are iteratively fed through an unmixing matrix, which is adjusted such that the information in its output channels is maximized. This minimizes the mutual information among them, thus rendering them independent of each other. ICA and its derivatives may be applied to a variety of problems. Perhaps its most successful application in neuroscience so far has been source separation in EEG data where ICA separated artifacts and various neural sources that contributed to the highly mixed EEG signal (Makeig et al., 1997b, 2002). With fMRI, ICA is primarily of use when there is uncertainty about the position and timing of activity, making its range of applications complementary to that of powerful hypothesis-driven methods such as statistical parametric mapping (SPM), which rely on temporal a priori knowledge (Friston, 1998). ICA has thus been shown to be useful in removal of artifacts and in the identification of nonpredicted but task-related activity (Calhoun et al., 2002; Duann et al., 2002; McKeown et al., 1998c). It should be noted that once the anatomical position of an area is known (through retinotopic mapping or functional markers), there is no need for data-driven analyses to visualize its BOLD response. Furthermore, the use of basis-functions allows SPM to work even without a priori hypotheses of the expected response shape (Buchel et al., 1998; Friston et al., 1995a; Henson et al., 2002). Compared with other data-driven methods, spatial ICA has been shown to perform best in isolating voxels whose activity time courses (ATCs) correlate with distinct task conditions or with

artifacts (McKeown et al., 1998c). Several applications and optimizations of ICA-fMRI have since been proposed (Calhoun et al., 2001a,b,c; Gu et al., 2001; McKeown, 2000b; Nakada et al., 2000; Nybakken et al., 2002; Stone et al., 2002; Suzuki et al., 2002). In our hands, ICA applied to conventional fMRI data sets of attention, emotion, and color-processing (Bartels and Zeki, 2000a,b; Zeki and Bartels, 1999) did equally well in identifying functionally activated areas as SPM, with the latter offering the advantage of sound statistical inference and a choice of a variety of statistical contrasts. There have been some exciting applications of ICA to fMRI data (Calhoun et al., 2002). However, it should be recognized that truly uncontrolled experiments are extremely rare, as the timing of stimulation or the onset of spontaneous brain activity is usually accessible to the experimenter even in experiments that do not fix such events a priori, for example, in the case of hallucinations or the reversal of visual illusions, thus allowing for an analysis using traditional statistical methods (Ffytche et al., 1998; Kleinschmidt et al., 1998; Lumer et al., 1998). Therefore, ICA applied to fMRI has, with some notable exceptions, been seen by many as a solution in search of a problem. This is likely to change with the combination of ICA with statistical inference, thus linking the strengths of both approaches (Calhoun et al., 2001b; McKeown, 2000a). Our application to free viewing data is, however, a model scenario for ICA: we assume that the brain's exposure to more natural viewing conditions will lead to distinct activity in a multitude of its functional subdivisions, and we have no a priori knowledge about the spatial location or about the activity waveforms in these activated regions (Bartels and Zeki, 1999, 2001), requiring us to use a data-driven method to segregate groups of voxels from each other whose ATCs differ. We hoped that ICA can identify the sets of voxels belonging to each distinctly activated functional subdivision that has an inherently distinct ATC and isolate them in a separate independent component (IC).

Stimuli and subjects

We analyzed two sets of fMRI data. All stimuli were projected onto a translucent screen of $26 \times 19^\circ$ visual angle, which was viewed through an angled mirror. For both imaging studies, informed consent was obtained from all subjects in accordance with the Declaration of Helsinki, and ethical approval was granted by the Ethics Committee of the National Hospital for Neurology and Neurosurgery, London, UK.

Study 1—epoch design

The first data set was an epoch-design study, which was analyzed by both ICA and SPM for comparison (see Fig. 1) and to validate ICA's performance on a large fMRI data set (200 whole-brain scans of $3 \times 3 \times 3$ mm resolution). In this first study, five normal subjects (three female, all right handed, aged between 24 and 38 years) viewed five types of stimuli arranged in repeated epochs of 30 s (Bork and Zeki, 1998). The stimuli consisted of recognizable and scrambled objects, which could be either moving or stationary, and a blank screen, thus allowing us to make the statistical contrasts shown in Fig. 1. The objects included cars, dogs, humans, bicycles, etc. Each epoch lasted 30 s and was repeated eight times in a fixed sequence. In each epoch, five stimuli of the same category were presented for 6 s each. Both static and dynamic images were derived from short video clips that had been processed by Adobe Photoshop and Matlab (Mathworks) (Bork and Zeki,

1998) to pixelate images such that they consisted of black squares (up to $0.6 \times 0.6^\circ$) on a white background (about 40% density).

Study 2—free viewing of a film

The second study provided data obtained in more natural free viewing conditions and was analyzed by ICA alone. In this study, eight volunteers (five female, all right-handed, aged between 24 and 38) viewed the first 22 min 25 s of the James Bond movie *Tomorrow Never Dies* (including the soundtrack) while blood oxygen level dependent (BOLD) activity was measured using fMRI (Bartels and Zeki, 2003, 2004b). The movie was interrupted every 2.5 or 3 min with a blank period (black screen, no sound) lasting 30 s, in total eight times, and for the purposes of another study the image was switched between gray and color every 30 s, which was, however, barely noticed by the subjects and not relevant to this study (Bartels and Zeki, 2003).

Acquisition of fMRI data

For both experiments, T2* weighted whole brain images ($3 \times 3 \times 3$ mm resolution) were acquired in a Siemens Vision 2 Tesla MRI scanner, using an echo planar imaging (EPI) sequence that optimized blood oxygen level dependent (BOLD) contrast. The following parameters apply to each study: For study 1 (epoch-design), 200 whole-brain images (64 slices; 2.5-mm thick with 0.5-mm gap) were acquired in a 20-min session using an echo time (TE) of 40 ms and a repetition time (TR) of 6.048 s. For study 2 (free viewing of a movie), whole brain images (48 slices, 1.8-mm thick and 1.2-mm gap) were acquired with a TE of 40 ms. For technical reasons, subjects 1–4 were acquired using a repetition time (TR) of 4.105 s with 324 whole-brain acquisitions in 22 min 12 s, subjects 5–8 with a TR of 3.649 s and 368 acquisitions in 22 min 23 s.

Data preprocessing and parametric analysis

All data were preprocessed using SPM99 (Friston et al., 1995b) (<http://www.fil.ion.ucl.ac.uk/spm/>) as follows. Whole-brain images were realigned to compensate for head movement and slices were temporally realigned to compensate for acquisition time lags. Images were spatially normalized to the Montreal Neurological Institute template (approximating to the atlas of Talairach and Tournoux, 1988) and spatially and temporally smoothed [with a Gaussian kernel of 6 mm full width at half maximum and by convolution with the hemodynamic response function (HRF) to reduce high-frequency noise]. After this preprocessing, both studies were analyzed by the data-driven method of independent component analysis (see below) (Bell and Sejnowski, 1995; McKeown et al., 1998c). To allow a comparison between the performance of SPM and ICA on a large fMRI data set, the epoch study (study 1) was also analyzed using SPM. Before SPM analysis, data were high-pass filtered with a cutoff of 300 s to reduce slow-frequency drifts. Then, a conventional multiple regression analysis was performed on each subject separately using SPM99, with one regressor for each of the five conditions, each consisting of a box-car convolved by the HRF.

Independent component analysis

The methodology we used for the ICA analysis in this and in our earlier studies (Bartels and Zeki, 2000a,b; Zeki and Bartels,

1999) was similar to the one of McKeown et al. (1998c), who first applied ICA to fMRI data. After preprocessing (above), voxels lying outside the brain were removed by manual intensity thresholding, done for each subject individually to reduce the total data to be processed. Data were then converted into a large matrix with the number of rows corresponding to the number of scans (200–368), and the number of columns corresponding to the number of voxels (between 60,000 and 70,000, depending on the subject). This matrix was submitted to the ‘runica’ procedure supplied in the EEG package by Makeig (<http://www.cnl.salk.edu/~scott/ica-download-form.html>) (Makeig et al., 1997a), which applied ICA incorporating the natural gradient feature of Amari (1999), using the default parameters. The total number of ICs obtained in each subject was equal to the number of input images, in our case ranging from 200 (study 1) to 368 (study 2); each IC has an associated time course, which corresponds closely to the BOLD signal of the most significant voxels in that IC. ICs shown in the figures were selected according to the correlation of their associated activity time courses: for study 1 (epoch design), the ICs whose associated time courses correlated most with any of the five stimulation conditions were selected (see Figs. 1 and 2). For study 2 (free viewing), ICs whose most significant voxels corresponded anatomically and which had significant intersubject correlations were selected (see below).

For display in the figures the relative voxel intensities (arbitrary units) in each IC were scaled such that the absolute of the maximum equaled 1 (most ICs had mainly neutral or positive voxels). Most figures show maximum intensity projections (MIPs), which were not thresholded, in sagittal and transverse views. MIPs display the voxels with the highest absolute value in the sightline (positive or negative).

Selection of ICs

One of the problems in the interpretation of fMRI-ICA results is the identification of ICs that contain cortical or subcortical functional units, and to separate them from potential artifacts, such as those induced by scanner noise, subject movement, breathing, or artery-pulsation, examples of which are shown in Fig. 4. Most previous applications of ICA could identify functional ICs by full or partial correlation with events or epochs (Calhoun et al., 2001a; Gu et al., 2001; McKeown, 2000a) or selected them according to both their temporal and spatial nature (Calhoun et al., 2002). However, in entirely uncontrolled experiments like the current one, a different method was needed as we had no spatial or temporal a priori hypotheses about the areas recovered by ICA. We predicted that functionally corresponding areas in different subjects should have temporally correlated ATCs, as identical stimuli should activate the same regions in the same way across subjects. We therefore screened ICs of each subject in four steps to select ICs containing potentially functionally specialized subdivisions. These included (i) a plausible anatomical location and distribution of the hottest voxels (e.g. ICs with the hottest voxels grouped together (such that they form an area), and lying in gray matter rather than e.g., in ventricles), (ii) bilaterality (e.g., ICs containing two groups of significant voxels in corresponding positions in both hemispheres), (iii) spatial correspondence across subjects (e.g., identification of ICs whose significant voxels form a similar shape in a corresponding position across several brains), (iv) intersubject correlation of activity time

courses (ATCs) (e.g., identification of ICs whose ATCs correlate across several brains, described below).

Once we had several sets of ICs whose most significant voxels corresponded anatomically across subjects, we used two methods to test whether their ATCs also corresponded, which would indicate that they share a similar function. A third, related technique was used to find out whether free viewing of a movie activated more cortical areas differentially than the epoch fMRI paradigm. In all analyses using ATCs, only periods of film viewing were used; the blank periods and the 30 s following them were excluded.

Testing candidate regions: intersubject correlations

To test whether anatomically defined candidate regions had significantly correlated ATCs, we calculated the pairwise correlation coefficients between the $n(n-1)/2$ intersubject pairs of IC-ATCs of this candidate region ($n = 8$ subjects in our study; see Table 1 for an example). A t test was then applied to determine whether these intersubject correlations differed significantly from 0. To rule out that these correlations were unspecific, i.e., could be accounted for by a correlation that was common to all candidate regions or even to all ICs, two controls were considered: first, pairwise correlation coefficients were calculated for ICs taken randomly from the pool of all candidate ICs (i.e., from each subject one out of the 10 shown in Fig. 3, which did or did not correspond anatomically to the others; ‘control 1’), or ICs were taken entirely randomly (one per subject, including candidate and noncandidate ICs; ‘control 2’). Calculations for both controls were iterated 1000 times.

Testing candidate regions: intersubject ranking

Calculating the mutual ranks (see Table 2) of candidate ICs (e.g., primary visual cortex in each subject) is a potentially more powerful way of determining how specific activity is to a set of anatomically corresponding areas than their mutual correlation coefficients. Ranking is relative to the remaining ICs within a subject, and therefore controls automatically for nonspecific correlations among a subject’s ICs (see below for an example). The rank is the position a candidate IC (e.g., the IC containing area V1, referred to below as ‘V1-IC’) has within all ICs of a subject when they are sorted in descending order by the correlation coefficient of their ATCs with that of another subject’s candidate IC (e.g., V1-IC of another subject). This way, given N subjects, one can calculate the rank that each of the remaining $(N-1)$ V1-ICs have concerning the V1-IC of the first subject. This can be repeated for the 2nd, 3rd, ..., n th subjects, leading to $N(N-1)$ ranks.

Table 1
Intersubject correlations of V1 ATCs

	A	B	C	D	E	F	G	H
A	1.00	0.38	0.43	0.62	0.51	0.65	0.63	0.38
B	0.38	1.00	0.19	0.19	0.24	0.35	0.33	0.25
C	0.43	0.19	1.00	0.31	0.25	0.43	0.36	0.32
D	0.62	0.19	0.31	1.00	0.45	0.48	0.64	0.37
E	0.51	0.24	0.25	0.45	1.00	0.35	0.50	0.28
F	0.65	0.35	0.43	0.48	0.35	1.00	0.57	0.32
G	0.63	0.33	0.36	0.64	0.50	0.57	1.00	0.29
H	0.38	0.25	0.32	0.37	0.28	0.32	0.29	1.00

Correlation coefficients between the ATCs of ICs that contain V1 between all eight subjects (A–H). Mean correlation coefficient: $r = 0.40 \pm 0.14$ SD; t test ($r > 0$): $P < 10^{-14}$, $n = 28$ (upper diagonal).

Table 2
Intersubject ranking of the ICs containing V1

	A	B	C	D	E	F	G	H
A	1	1	1	1	1	1	1	1
B	2	1	5	6	3	2	2	1
C	2	9	1	3	6	2	2	1
D	1	4	1	1	1	1	1	1
E	1	2	2	1	1	3	1	2
F	1	1	1	1	2	1	1	1
G	1	1	1	1	1	1	1	1
H	1	1	1	1	4	1	1	1

Ranks of V1-ICs when all ICs of one subject are sorted by their ATCs’ correlation with other subjects’ V1-ICs. On the example of the first row, the numbers indicate the rank of V1-IC of subject A when all its ICs are sorted in descending order by their ATCs’ correlation coefficient with the V1-ICs of subjects A to H: each time, A’s V1-IC ranked first. When C’s V1-IC was ranked with B’s V1-IC, it came on 9th position. Random intersubject correlations would lead to an even distribution of the ranks between 1 and N , with N = number of ICs in a subject (here, between 324 and 368). Here, the median rank equaled 1 [expected: $173 (= (324 + 368) / 2)$], Kolmogorov–Smirnov one-sample test: ranks are higher than expected, $P < 10^{-47}$, n (ranks) = $8(8-1) = 56$]. Note that for later applications, ranks were normalized to values from 0 to 100.

These are twice as many as pairwise correlation coefficients, as ranks are not symmetrical: if V1 of subject C appears in rank 6 when its ICs are sorted by subject E’s V1, the latter might appear in, e.g., rank 2, when its ICs are sorted by the former (see Table 2). We normalize ranks to 0–100 to become independent of the number of ICs per subject, and we report the median rank rather than the mean rank, because the distribution of ranks is not normal. A Kolmogorov–Smirnov test (not dependent on the normal distribution) can be used to establish whether the distribution of ranks obtained differs from that expected for ATCs of nonspecific correlation, which would be evenly distributed between 0 and 100. An advantage of intersubject ranking compared to correlation is that it provides a relative measure: if many areas in a subject have highly correlated ATCs but no area-specific activity, intersubject correlations can be high, but ranks would still be low. We use the ranking-system to reveal the specificity of activity in a given area with respect to others: for example, V1 correlates best with V1 in other subjects, but worse with other candidate regions (see Fig. 6). The same controls as for intersubject correlation were applied (i.e., a within-candidate regions control and an overall control).

Comparison of different studies: intersubject correlations

In context of our aim to dissect the brain into functional subdivisions, it is important to determine which stimulus can activate more regions and more differentially. The distribution of intersubject correlation coefficients in a given study is a useful indicator for this, as it reveals the number of areas (or ICs) that have similar ATCs across different subjects. In other words, this measure can be used to compare the effectiveness of different stimuli (e.g., movie viewing vs. a traditional epoch paradigm) in terms of the number of distinct areas they activate differentially. The distribution of intersubject correlation coefficients (i.e., of IC-ATCs between all pairs of all subjects) has $[N(N-1)/2] \times n_{ics}^2$ intersubject correlation coefficients, with N being the number of subjects and n_{ics} the number of ICs per subject (see inset in Fig. 9). If obtained for two separate studies with identical number of ICs per subject, a two-sample Kolmogorov–Smirnov test (which does not assume normal distribution) can be applied to determine

the statistical difference between the distributions. Intersubject correlation distributions are however dominated by the high number of low correlations. As a measure that is more sensitive for revealing the quantity of ICs with high intersubject correlations we used the measure of *maximal intersubject correlations*. We obtained this by identifying for a given IC the IC in each of the other subjects that correlated best with it and to take for each IC the average of these best intersubject correlations. For each IC (of each subject), one obtains therefore one number, which is the average of the best correlations this IC has with the best-correlating ICs of the other subjects, a total of $N \times n_{\text{ics}}$ *maximal intersubject correlations*. The distribution of these can then be plotted for different studies (see Fig. 9). Both measures (intersubject correlations and maximal intersubject correlations) are independent of the number of subjects, but depend on the number of ICs per subject. For our comparison of the free-viewing study with the epoch study, we therefore reanalyzed the former using a limited number of scans to match the number of ICs obtained with that of the epoch study. As a control and to obtain a baseline for the distributions expected for completely random ATCs, we generated an artificial data set of eight subjects each with 200 ATCs that were obtained by convolving white noise with the hemodynamic response function. We note that the danger of a ‘super-fractionation’; that is, when ICA artifactually breaks a single area into several components, which could be a potential problem in comparing different studies would, if at all, lessen the difference between the distributions obtained from the two experimental data sets: the large, unspecifically activated networks of areas obtained with a simple stimulation paradigm are more likely to be split up into smaller fragments than the individual areas obtained with a complex stimulation, and lead to higher intersubject correlations. In addition, it is a conservative measure for the number of specifically activated areas because it only accounts for areas that are highly stimulus driven, while ignoring those with more complex and less stimulus-driven activations, such as those dealing with higher cognitive functions, subject-specific associations and emotions.

Results and discussion

Experimental comparison of SPM and ICA

To convince ourselves that ICA delivers physiologically meaningful results, we began by analyzing data from a more traditional epoch-design study intended for parametric statistical analysis, in which five subjects had viewed five types of stimuli arranged in repeated epochs of 30 s. The stimuli consisted of recognizable and scrambled objects, which could be either moving or stationary, and a blank screen (Bork and Zeki, 1998). An SPM analysis (Friston et al., 1995b) showed that all nonblank stimuli activated the early visual areas V1/V2 (Fig. 1d), stimuli containing motion activated the V5 complex (Fig. 1e), while stimuli containing objects activated the object and face selective subdivisions of the lateral fusiform gyrus in the lateral occipital cortex (Fig. 1f). The latter is likely to contain several subdivisions, and is also known as lateral occipital complex (LOC) (Kanwisher et al., 1997; Malach et al., 1995). We were now curious to see what a purely data-driven analysis, without any a priori knowledge of the number of epochs, their onsets, etc. would make of the same data. If ICA is able to segregate regions that differ in their ATCs, one would expect the

same spatially segregated regions that were shown by SPM to differ in their ATCs to appear, each in a separate IC. This turned out to be so. ICA revealed the same subdivisions as SPM in the occipital lobe, each in a separate independent component (IC) (Figs. 1a–c), thus showing that ICA can segregate regions that are differentially activated by stimulus epochs. Each of these ICs had a time course that closely matched the BOLD time course of the hottest voxel revealed by SPM in the corresponding region (see Figs. 1a–f). The results were highly reproducible across subjects. This is shown in Fig. 2, where the similarity of activation revealed by ICA across four subjects is illustrated. For each subject, the ICs containing early visual areas (V1/V2) and motion-selective areas (V5) are shown along with their associated ATCs. These ICs were the ones whose ATCs correlated most with the experimental task conditions. Both the spatial extent and the associated ATCs were highly conserved across different brains, which demonstrates ICA’s reliability in recovering physiologically meaningful ICs. More importantly, we note that the ATCs associated to the ICs of V1/V2 and V5 have each distinct and specific waveforms that are area-specific and preserved across subjects. Note that the ATCs not only reveal a distinct preference of each area for distinct epochs, but that specificity is even preserved within epochs. V1 for example shows a pronounced recovery during the rest period, while the signal in V5 remains constantly low during rest. Such distinct area-specific temporal characteristics have been observed before (see e.g. d’Avossa et al., 2003) and are encouraging for our hypothesis that distinct areas should be segregatable by their unique ATC fingerprints. At the same time, this observation should be a caveat to the many SPM users who use standard box-car regressors in SPM—these might coincidentally fit better to the BOLD response of one area than of another, and thus favor ‘activation’ of some regions compared to others (see also Methods). In conclusion, we found that (i) ICA performed as well as SPM in identifying areas that were activated differentially, even without a priori knowledge, (ii) both the activation of distinct areas and ICA’s performance across subjects was highly reliable and reproducible, (iii) the ATCs associated to each IC matched closely the BOLD signal of the ICs hottest voxel, (iv) functionally distinct areas exhibited differential ATCs, not only because of differential stimulus preference, but also due to distinct area-specific dynamics. In the following study, we wanted to find out whether these factors can be used to segregate and map a multitude of distinct areas based on their differential temporal fingerprints.

Functional brain anatomy revealed through natural viewing conditions

The above results encouraged us to use ICA in a domain inaccessible to parametric statistics, namely the ‘blind’ mapping of functional subdivisions in the brain. To optimize our chances to succeed, we had to adopt an approach virtually opposed to that used in conventional fMRI, both in terms of its aim, the experimental procedure and the analysis. Traditional experiments typically aim to differentially activate as few areas as possible—optimally only one—to determine an area’s functional role. Highly specific and often very abstract stimuli are used to achieve such selective activation, and the analysis is based on a priori knowledge of the stimulus and of the expected response to explicitly test the experimenter’s hypothesis. Our approach was entirely different. In contrast to traditional experiments, our aim was to activate as many regions as possible, each with a differential ATC. According to the

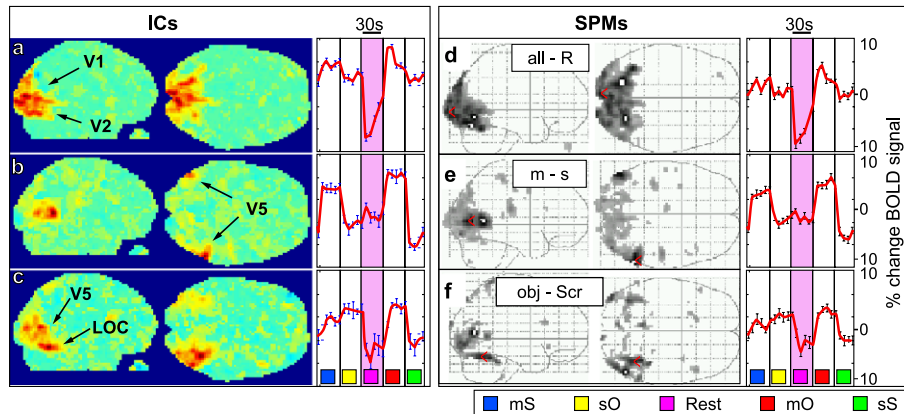


Fig. 1. Comparison of an SPM and an ICA analysis of single subject data from a traditional epoch design fMRI study on motion and object recognition. (a–c) The three ICs obtained by ICA whose ATCs correlated most with the stimulus conditions. They revealed differential involvement of primary visual (a), visual motion selective (b) and object selective areas (c) in different stimulus conditions, as is apparent in the ATCs shown to their right. ICs are shown as maximum intensity projections and were not thresholded. The color-code of relative voxel intensities is shown in Figs. 4 and 5. ATCs of the ICs (a–c) have arbitrary units and were averaged over the eight repeats of the conditions. (d–f) Statistical parametric maps (SPMs; $P < 0.001$, uncorrected) show the same cortical areas as revealed by ICA. BOLD signals taken from the most significant voxel are shown to the right of each SPM (averaged over condition repeats). Error bars: SEM Abbreviations: mS = moving scramble, sO = static object, mO = moving object, sS = static scramble. [all–R] = all stimuli vs. rest, [m–s] = motion vs. static, [obj–scr] = objects vs. scramble.

principle of functional independence, this would be best achieved by exposing the brain to natural conditions—these are complex, perceptually rich, and have multiple attributes that vary independently (Bartels and Zeki, 2003, 2004b). Distinct functional subdivisions of the brain should therefore differ maximally in terms of their activity time courses. A data-driven analysis such as ICA should then be able to segregate voxels belonging to functionally specialized areas based on their distinct ATCs, thus leading to a time-based dissection of the brain. In contrast to the traditional approach, such a stimulation would be entirely uncontrolled and nonrepetitive, and no a priori knowledge about stimuli or response would be needed for the analysis, apart from the general expectation that natural conditions would be optimal to elicit differential

temporal fingerprints in distinct areas. To approximate a stimulation of the brain with more natural conditions, we asked eight subjects to freely view the first 22 min of the James Bond movie *Tomorrow Never Dies*. This was a first and undoubtedly crude attempt to map the brains' chronoarchitecture. There were two factors that could potentially foil our approach. First, the brain was exposed to rapid, real-life stimuli, while we measured the slow BOLD response that has time constants in the range of several seconds, sampled at intervals of 3–4 s. There was therefore a danger that all temporal variability of neural activity was smoothed out by the slow BOLD response, leading to a uniform and indistinguishable signal in all stimulated areas. Second, one would naturally expect from hypotheses such as those of Schiller (1997), which suppose that perceptual




Fig. 2. Consistency of ICA results across four subjects in the epoch-design experiment on motion and object recognition. For each subject, the ICs containing V1 (left) and V5 (right) are shown. These ICs had ATCs that correlated most with the task conditions. To the right of each IC, its associated ATC is displayed, averaged over the eight stimulus repetitions (peristimulus time histograms; PSTHs). Note that ATCs are specific to the area, and similar across subjects. For abbreviations, see Fig. 1.

analysis “. . . is performed interactively by areas and neurons with multipurpose properties” that, especially in natural conditions, the temporal characteristics of different areas would be barely distinguishable, thus making it difficult to separate them on this basis. Our results show that both factors did not apply—functionally distinct regions had distinct ATCs during free viewing.

Independent components contain distinct cortical subdivisions

The ICA analysis revealed a multitude of areas in the visual and nonvisual cortex in every subject, some of which corresponded from subject to subject in terms of their anatomical locations and shape, while others were subject-specific. Most of these areas were contained in a separate IC each, which therefore contained only one hot spot of voxels, or two hot spots on bilaterally corresponding sites. Some of the ICs contained in addition to their most significant cluster/area additional clusters of comparably more faintly activated regions, thereby indicating a high degree of temporal correlation and possibly anatomical connectivity between them (Bartels and Zeki, 2004a). In almost every subject, we identified 10 ICs whose most significant voxels corresponded across subjects with a high degree of consistency. Each of these 10 ICs contained a separate cortical region (most were bilateral), and was identified either in every subject or in the majority of subjects, leading to a total of 71 ICs across the eight subjects. We concentrate in the following analyses on these 10 ICs and their ATCs, as their anatomical correspondence across subjects makes it likely that they reveal functional subdivisions that reflect a general property of the human brain. The 10 ICs / areas that were identified most consistently in each subject are shown in Fig. 3, on the example of a single subject. Each IC was intensity-thresholded (allowing only voxels with a positive activation of >30% to pass) and color-coded before it was superimposed on the individual's structural image. For simplicity, we refer to the cluster of the most significant voxels in a given IC as to a separate ‘area’ below, and emphasize again that each ‘area’ was isolated in a separate IC in each subject (with the exception of a region in the precuneus and the retrosplenium, which were isolated together in a single IC). Furthermore, we note that, at this stage of the analysis, it was still undetermined whether these ‘areas’ truly reflected cortical subdivisions with distinct functions. The anatomical locations of the 10 areas / ICs correspond to the following regions (with the probable functional subdivision in brackets): (i) the calcarine sulcus [V1], (ii) the anterior calcarine [presumably the peripheral field representation of V1/V2], (iii) the posterior fusiform gyrus [ventral V2/V3], (iv) the middle fusiform [the color-sensitive V4 complex], (v) the posterior middle temporal cortex [the motion-sensitive area V5], (vi) a posterior [LOp] and (vii) a lateral [LOl] subdivision of the object and face selective region in the fusiform gyrus [also known as lateral occipital complex (LOC) and including FFA; Kanwisher et al., 1997; Malach et al., 1995], (viii) a parietal area in the precuneus, with a second hot spot in the retrosplenium, (ix) the auditory cortex [BA 41 / 42] and (x) a region corresponding to Wernicke's receptive speech area [Brodmann area (BA) 22]. The fact that a considerable number of ICs contained regions that had a clear anatomical correspondence between many or all of the subjects was a first indication (i) that the functional organization of the brain during natural conditions was preserved across subjects, (ii) that ICA detected the voxels that belong to each of these functional subdivisions reliably in every subject by their characteristic ATCs, and (iii) that cortical subdivisions were segregated in a separate IC each. Before we set out to prove that

each of these ICs truly reflects functional subdivisions of the brain based on their differential ATCs, we briefly address their counterparts: ‘artifactual’ ICs.

Artifacts

In addition to above ‘functional’ ICs, every ICA produces also ‘artifactual’ ICs. Typical examples of the most commonly encountered artifactual ICs are shown in Fig. 4. These include components containing non-neuronal structures such as eyes (a), ventricles (b) and blood vessels (not shown), scanner-induced signals such as spikes (c), movement artifacts involving different parts of the surface of the brain (d), and ‘noisy’ ICs (e), the latter making up about a third of ICs in a typical data set. We note in this context that most of these ‘artifactual’ ICs represent in fact a true signal, though one in which we are not interested here, or one whose functional significance is difficult to prove. In some cases, we observed that functional areas, such as V1, were split into subdivisions, with one component containing the main part of V1 and another, for example, a dorsolateral part. Such ‘superfractionations’ may appear because ICA is forced to produce a larger number of ICs than is actually present in the data. They are, however, easily detectable as they are not consistent across subjects and because fractions belonging to the same functional area have very high mutual correlations between their ATCs. In our data, most of the candidate ICs shown above were indistinguishable when the same data were analyzed as to retrieve 368 or 200 ICs (see below) showing that ICA did only rarely ‘superfractionate’ actual functional components. In addition, there were numerous ICs that contained clean, often bilateral clusters of voxels in gray matter, whose spatial distribution was, however, either not repeated across subjects or whose ATCs showed no relation across subjects. These ICs contained most likely specialized regions whose activity was not task-related, and of higher cognitive, emotive or mnemonic nature. Because of the difficulty to confirm the functional significance of such ICs we concentrate in this report on ICs whose functional significance is task-driven and therefore verifiable. Alternative ways of relating non-task-driven ICs to brain function are discussed briefly toward the end of this report.

Distinct ICs have independent, stimulus-driven and area-specific ATCs

Having found ICs that showed anatomically corresponding regions across subjects, we expected two things: first, that ICs whose regions corresponded across subjects would have similar ATCs between subjects, since these ICs were likely to represent areas of identical functional specialization that were similarly stimulus-driven in each of the volunteers. Second, we expected anatomically non-corresponding ICs to have low correlations, since ICs that differ in their anatomical location were expected to differ also in their function. Both expectations were confirmed by our data. ICs that contained anatomically corresponding areas in different subjects had time courses that correlated, indicating a functional correspondence: the same stimulus activated distinct areas within each brain in a distinct way, and consistently across separate brains. An empirical demonstration of this is given in Fig. 5 where ATCs of 4 of the 10 anatomically corresponding ICs, arbitrarily taken from the first three subjects, are superimposed. The ICs shown contain




Fig. 3. Ten chronoarchitecturally identified areas (each in a different IC) of a single subject who freely viewed the movie *Tomorrow Never Dies*. Each region was identified by ICA in a separate IC, which was then color-coded, intensity-thresholded (>30% positive activation) and superimposed onto the subjects' brain. Each area (IC) had an area-specific ATC that correlated significantly and exclusively with anatomically corresponding areas (ICs) in the other eight subjects (see following figures). ACS = ventral lip of the anterior calcarine sulcus, Aud = auditory cortex (BA41 and 42), Lol = lateral part of the lateral occipital complex (LOC), Lop = posterior part of LOC, pc + rs = network containing precuneus (BA7) and retrosplenium (BA23 or 30), Wern = Wernicke's area (BA22).

cortical areas corresponding to the visual cortical areas V1, V4, V5 and the auditory cortex. A visual inspection of the superimposed ATCs suggests that each of these areas has a time course that is characteristic of it and preserved across subjects. This is quantified and statistically verified in Fig. 6 for the 10 cortical ICs / areas (71 across all subjects; shown in Fig. 3) that were identified most consistently in each subject. Fig. 6 shows for each of these 10 ICs its mean intersubject correlation and its median intersubject rank (see Methods). We found that each of the ICs/areas had significantly correlated ATCs with its anatomically corresponding counterparts across subjects [each with $P < 0.0001$, t test on $n(n-1)/2$ correlation coefficients, with n being the number of subjects; see Methods and Table 1 for an example]. This indicates that these areas were (i) stimulus-driven and (ii) processing corresponding features of the movie. In average, the between-subject correlation of ATCs of anatomically corresponding ICs was $r = 0.30 \pm 0.16$ SD ($n = 222$). The controls show that this significant intersubject correlation was abolished when the compared areas did not correspond anatomically across subjects (control 1: $r = 0.03 \pm 0.16$, $P = 0.27$; for each subject one of the above 10 areas was randomly selected; control 2: $r = 0.00 \pm 0.08$, $P = 0.50$; for each subject a random IC from all ICs was selected; for both controls, the results of 1000 combinations were averaged). Similar results were obtained for intersubject ranking. When the ATC of one subject's IC was used to rank-order all ICs of the remaining subjects, the first ranks were occupied by ICs that corresponded anatomically to this initial IC (see Methods and Table 2 for an example). In average, intersubject ranking led anatomically corresponding ICs to appear in the median rank of 0.54 (1st quartile: 0.27; 3rd quartile: 2.88; ranks normalized to values from 0 to 100). Kolmogorov–Smirnov tests showed that the distribution of rank orders were skewed toward the top ranks for each of the 10 cortical areas [each with $P < 10^{-13}$, KS-test on $n(n-1)$ ranks]. The controls show that these findings cannot be accounted for by unspecific correlation of ATCs, and that they must be due to area-specific processing that is similar in each subject (control 1: median rank: 44, quartiles: 11, 78, KS test: $P = 0.11$; control 2: median rank: 50, quartiles: 26, 74, KS test: $P = 0.43$).

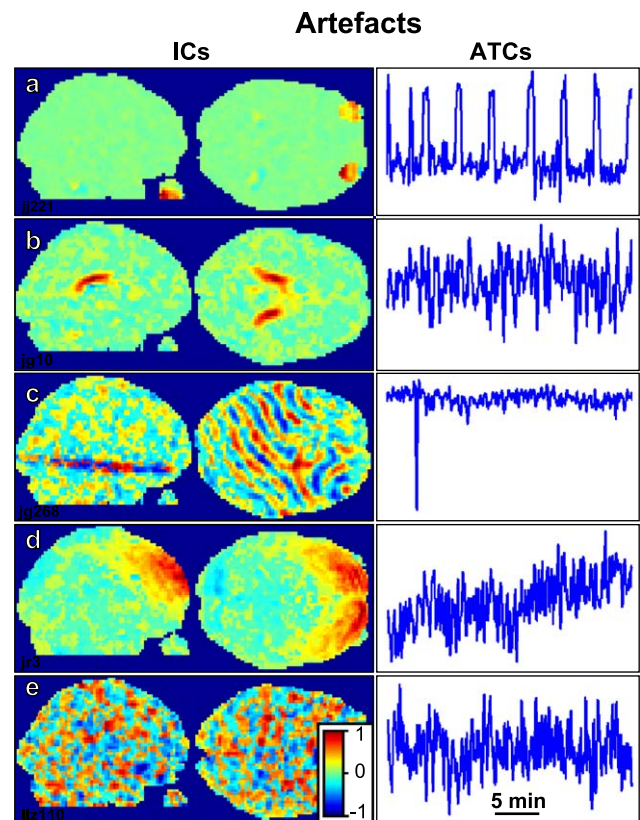


Fig. 4. Representative artifacts together with their ATCs. (a) Eyes. ATCs reflect eye movements, which are reduced during the eight blank periods that interrupted the movie. (b) Part of the ventricles. In each subject, the complete ventricles were isolated, but fractionated into separate ICs, probably due to fluid flow. (c) Scanner-induced 'spike', affecting only one slice at one point in time. (d) Movement artifact, affecting a large region on the brains' surface, with a steadily increasing signal. (e) 'Noise' artifact. About 30% of ICs in an analysis contain this type of noise.

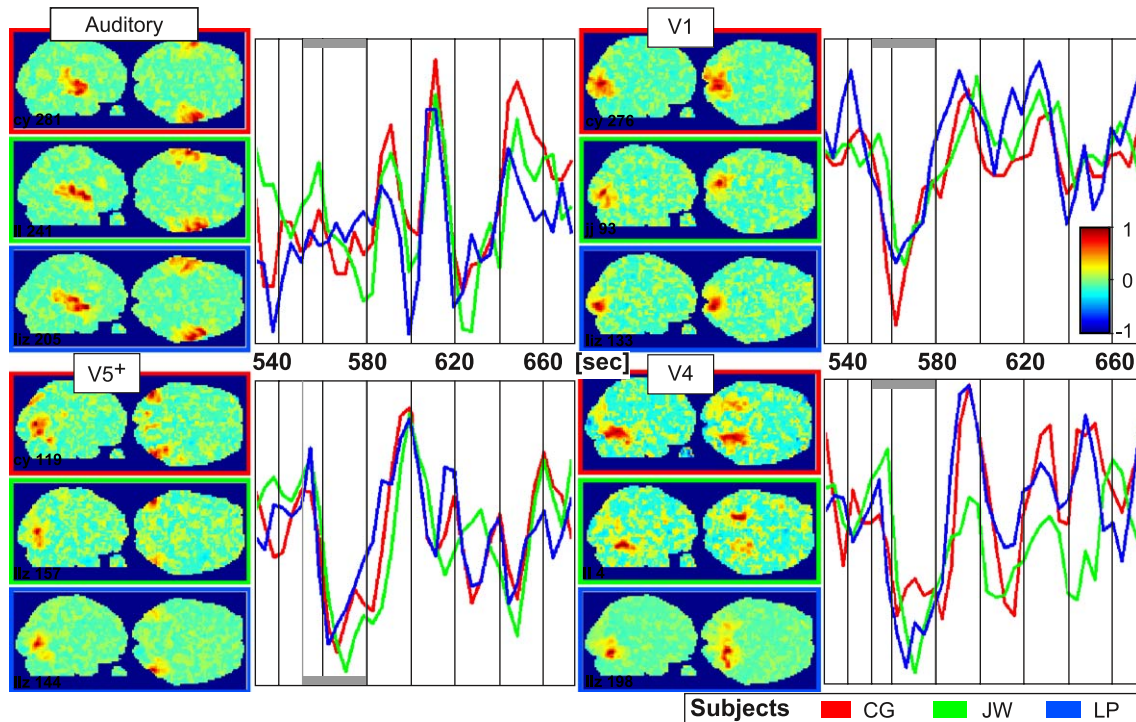


Fig. 5. Anatomically corresponding areas in different brains had similar ATCs during free viewing of the movie. For four cortical regions (auditory cortex, V1, V5 and V4) the ICs of the first three participants in our study are shown, along with their superimposed ATCs (arbitrary units on the y scale). The gray bars at the top and bottom indicate blank periods (black screen, no sound), which were not considered in further analyses (see Methods).

We can draw the following conclusions from above results: First, the fact that anatomically corresponding areas in different subjects had significant ATC correlations demonstrates that these candidate regions were stimulus-driven. Second, the lack of corre-

lation between anatomically noncorresponding regions demonstrates that distinct regions responded differently to the same stimulus, therefore proving that anatomically noncorresponding regions had distinct functional specializations. These results there-

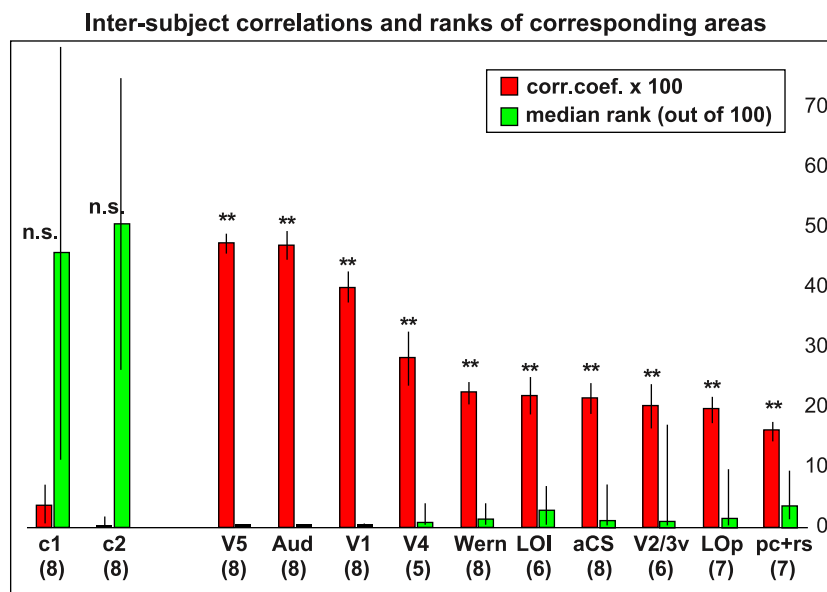


Fig. 6. Between-subject correlations of ATCs were specific to anatomically corresponding areas. For each of the 10 areas shown in Fig. 3, the mean intersubject correlation coefficients ($r \pm \text{SEM}$) and the median intersubject ranks (rank \pm quartiles, normalized to 100) are shown. When, instead of anatomically corresponding areas, random areas out of the 10 (control 1, c1) or any random IC (control 2, c2) were chosen, correlations were no longer different from zero, and ranks not skewed towards zero. $**P < 0.0001$ (t test on correlation coefficients); $P < 10^{-13}$ (Kolmogorov–Smirnov test on ranks). See Fig. 3 for abbreviations. n = number of subjects with that area.

fore demonstrate that the human brain can be dissected into functionally distinct subdivisions based entirely on activity time courses recorded during natural viewing, with a high reliability across subjects.

The independence of ATCs within brains and the principle of functional independence

Chronoarchitectonic mapping works in the first place because distinct areas appear to have distinct ATCs during free viewing. The reason for this temporal independence of distinct areas is probably a combination of both, the high degree of temporal independence of the distinct features that are processed in each area, as stated by the principle of functional independence, and the distinct physiological, connectional and molecular properties of each area, which lie at the basis of their functional specialization.

To obtain an estimate for the degree of independence between distinct (anatomically noncorresponding) areas, or, in other words, to obtain an estimate for the degree of communication between areas within a single brain, we calculated the mean correlation coefficient of ATCs of above 10 candidate ICs (per subject) within each of the eight brains. In Fig. 7, this is contrasted with the mean correlation coefficient of anatomically corresponding regions between subjects and with that of anatomically noncorresponding regions between subjects (all using the same 71 ICs, in different pairings). Anatomically corresponding areas between different subjects had a correlation of $r = 0.30 \pm 0.16$ SD, $n = 222$ (same for absolute values). The average within-brain correlation between (anatomically distinct) areas was $r = 0.12 \pm 0.21$ SD, $n = 281$ (absolute values: $r = 0.19 \pm 0.14$), and the average correlation between anatomically noncorresponding areas between subjects was $r = 0.00 \pm 0.08$ SD, $n = 2520$ (absolute values: $r = 0.06 \pm 0.05$).

It is a striking testimony of the independence of distinct areas during free viewing that anatomically corresponding areas between different subjects had a much higher correlation than the average within-brain correlation between these areas ($P < 10^{-12}$ for both, a

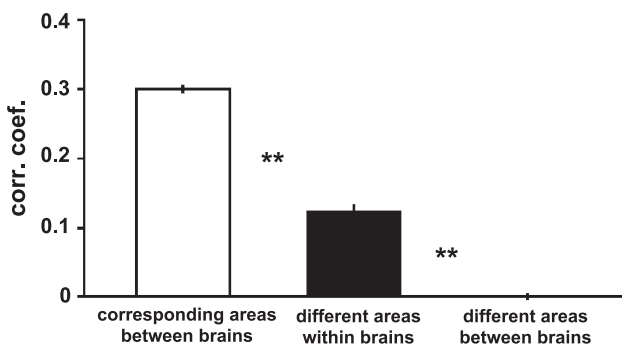


Fig. 7. ATC correlations between and within brains indicate a high degree of independence between distinct areas, even within a single brain. The high correlation of anatomically corresponding areas between brains indicates their functional correspondence. In contrast, the low ATC correlation of distinct areas within the same brain reveals a high degree of temporal and therefore also functional independence during natural viewing. There is no correlation between noncorresponding areas between brains. Data include all corresponding pairwise correlations of the 71 stimulus-driven ICs shown in Figs. 3 and 6, across or within all eight subjects (see also text). Wilcoxon rank sum test and two-sample t test: $P < e^{-12}$ between each set ($n = 222, 281, 2520$). Error bars: SEM.

two-sample t test and a distribution-independent nonparametric Wilcoxon rank-sign test). The comparably higher correlation of anatomically noncorresponding areas within the same brain compared to those between different brains can be accounted for by both, a physiological factor and a methodological confound, which are difficult to segregate: the communication of areas within a brain will partly account for an increased within-brain correlation, while the remainder will originate from artifacts that can affect every voxel of a given brain, such as breathing and head movements. Whatever their relative contribution may be, the low within-brain correlation highlights the surprisingly high degree of independence between distinct cortical areas during free viewing.

If each area was driven by distinct, independently varying features in the movie, it must have been their lower-frequency components, together with area-specific properties, that allowed the high degree of temporal independence to emerge between areas despite the slow BOLD response and the low sampling frequency. High-frequency variation in the stimulus and in the neural responses would have been smoothed out by the slow hemodynamic response, which has time constants of several seconds. Fig. 8 shows the average power spectrum of the ATCs of the 71 stimulus-driven ICs from all subjects, together with the power spectrum of artificial ATCs derived from white noise convolved with the estimated HRF as it is provided in SPM99. While there is some power in the range of seconds, the first peak is reached at a period of about 30–40 s, which is about 10 s more than expected based on the artificial ATCs. This indicates that natural stimuli themselves have considerable power at low frequencies, at least in their interaction with the physiological properties of cortical regions. The high power at very low frequencies is likely derived from a combination of slow physiological drifts, but also head-movements and scanner-induced drifts.

Our results of a high temporal independence of functional subdivisions during natural conditions match closely those of our previous analysis of natural viewing data that included brain regions that were identified functionally, not data-driven like here (Bartels and Zeki, 2003). This previous study additionally showed that the low correlation between brain regions was directly related to the low correlation of features (such as color, faces, etc) as subjectively experienced during free viewing. At the same time, it showed that activity in specialized regions correlated linearly with the intensity of the subjective experience of the attribute that region is specialized for, thus providing experimental evidence for the principle of functional independence (Bartels and Zeki, 2003).

We should, however, emphasize that the functional independence between distinct areas is ultimately derived from structural factors, which themselves will have contributed to the observed differential BOLD responses. These factors may not only include distinct local connectivity, synaptic composition, etc., but also the density of vascularization, which can be specific to different functional cortical subdivisions (Cavaglia et al., 2001; Harrison et al., 2002; Zheng et al., 1991). If these structural factors played a dominant role, one would of course expect the complexity of the stimulation to play a comparably smaller role in eliciting differential responses in distinct areas. We were therefore curious to find out whether a more complex stimulation leads to more differential responses in more areas, as would be predicted by the principle of functional independence. In the next section, we compare the efficiency of natural free viewing conditions with that of conventional fMRI paradigms in activating more areas differentially.

Free viewing activates more areas differentially than traditional fMRI paradigms

We hypothesized above that, according to the principle of functional independence, exposure to natural conditions should activate more brain areas with more distinct ATCs than exposure to more simple stimuli. To test this, we compared the data obtained from our free viewing study with those obtained from the more classical object-motion study. This was a conservative comparison since, compared to other traditional fMRI studies, our epoch study used a complex and perceptually rich stimulation; its ‘moving objects’ stimuli were in fact 6-s movie clips, showing cars, people or animals driving, cycling or walking along streets, fields and other complex scenes. Moreover, subjects were not constrained in their eye movements (Bork and Zeki, 1998). Our comparison therefore amounted to ask whether more complex stimuli activate more areas more differentially than less complex stimuli. To answer this, we used two measures: intersubject correlations and maximal intersubject correlations (see Methods), both of which indicate how many ICs have ATCs that correlate across subjects and therefore reveal how many stimulus-driven regions were differentially activated in each study. In an equivalent comparison that controlled for the number of scans (and thereby also the number of ICs), TR and number of subjects (see Methods), we obtained dramatic differences in the number of differentially activated regions in the two studies: when we counted the average number of ICs per subject pair that had an intersubject correlation of $r > 0.5$ (which is a measure for the number of differentially activated, stimulus-driven regions), we obtained 66 ICs in the free-viewing study, compared to only 4 ICs in the epoch study, and 0 in the control. A two-sample Kolmogorov–Smirnov test showed that the intersubject correlation distributions obtained from the free-viewing and from the epoch study differed significantly ($P < 10^{-6}$), while those of the epoch study and the control did not differ ($P < 0.44$). The highly increased number of ICs with intersubject correlation in the free-viewing study can only be explained by the fact that the more complex stimulus activated many more areas with different ATCs, and each in a distinct way, though consistently across subjects. The full distributions of intersubject correlations and of maximal intersubject correlations are given in Fig. 9. These distributions reflect the effectiveness of distinct stimuli to differentially activate different areas, and can be used to compare different data sets that were collected under different stimulation conditions but with the same analytical parameters. The above finding does not demonstrate that more areas are active during more complex tasks, but that more areas have differential and characteristic ATCs during such tasks. This differential (or more independent) activity in more areas during exposure to natural conditions is important for several reasons. First, it shows that area-specific ATCs can be obtained despite stimuli that are uncontrolled and whose dynamics are much more rapid than the measured BOLD signal. Second, it demonstrates that areas maintain a high degree of functional specificity in natural conditions, which is contrary to the view that many ‘multipurpose’ areas process many features unspecifically in large cooperating networks (Schiller, 1996, 1997). Third, it is a necessary step in the experimental proof of the principle of functional independence outlined above, in that more areas obtain more characteristic ATCs in more natural conditions, which is most likely accounted for by the independence of features that are processed by the brain in separate areas.

Two steps to determine the function of stimulus-driven ICs

It is a particular strength of the chronoarchitectonic approach that, like other anatomical techniques, it provides a functional map of the brain that is independent of experimenter’s hypotheses, correctly controlled stimuli and human interpretation. This makes it more likely to provide a more objective map and reveal unexpected and novel subdivisions. However, the benefit of objectivity comes at the cost of ignorance concerning the function of the regions identified. Although identifying a function of a region is far more demanding and fraught with human error than identifying it as a functionally separate entity, we discuss in this and the following section two separate, mutually complementing procedures to gain information about the potential function of a novel region. Since our previous report addresses part the procedure in detail, and since we will examine particular regions identified here in future studies, we concentrate on concepts here.

The most obvious approach to identify the function of an area in natural viewing experiments would be to relate the activation of the region to the stimulus, in the sense of a reverse correlation (Hasson et al., 2003). Concretely, one would for example specifically examine those sections of the stimulus (in this case, short sections of the movie) when the area of interest is at relative peaks of activation, and compare them to sections when the area is at relative lows. We note, however, that this is a highly empirical procedure that leaves much room for human interpretation. Especially for regions involved in attention, emotion or cognitive aspects, one has probably to rely on human intuition rather than on low-level features of the stimulus, and it may be difficult to come up with the correct interpretation. We see this therefore not as a suitable formal test for the functional identification of a region, but more as a useful way to generate hypotheses about its function. We therefore propose a two-stage procedure: in the first stage, one uses reverse correlation to generate hypotheses about possible functions, which are then formally tested in the second stage, as we have demonstrated in our previous report (Bartels and Zeki, 2003). For example, in the first stage, one hypothesis is that a region is involved in face recognition, while the other says it is involved in the processing of whole human bodies. In the second stage, one can then obtain a measure that indicates the intensity with which the features in question vary over time in the stimulus, for example, by asking human observers to rate the intensity with which they perceive these features continuously throughout the movie. The obtained attribute-specific intensity time courses can then be formally regressed against the ATC of the region in question, or against the whole brain activity, as it is done in traditional statistical parametric mapping (SPM) (Bartels and Zeki, 2003). This provides a formal test about which of the hypotheses is more valid. It will also show how much of the variance in a region is explained by the variation in the stimulus feature. And if a whole-brain regression is performed, the result will show which other regions in the brain correlate significantly with the intensity of the feature in question (Bartels and Zeki, 2003). We have demonstrated that even a regression of the intensity of a single feature against natural viewing data can lead to highly area-specific results, and that the perceived feature intensity correlates linearly with activity in regions that are known to be specialized for that feature (Bartels and Zeki, 2003). This makes it likely that a correct hypothesis has much chance of succeeding to identify the specialized regions in question.

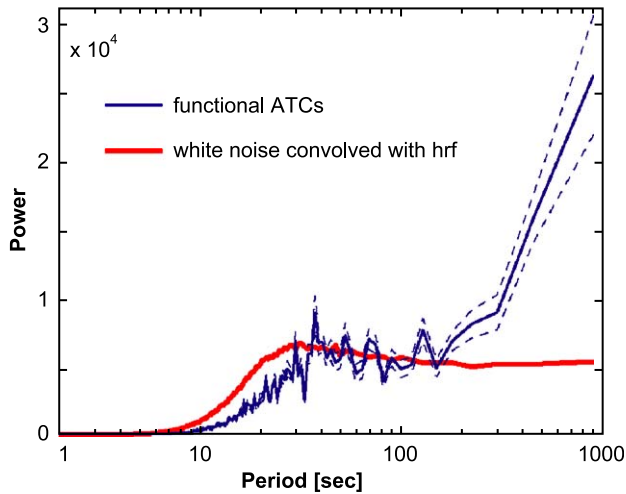


Fig. 8. Power spectra of ATCs of stimulus-driven ICs (blue) compared to those of artificial ATCs consisting of white noise convolved with the hrf (red). Power spectra were calculated for each of the ATCs of the stimulus-driven ICs considered in previous figures (71 ICs from all subjects) and then averaged. The broken line indicates SEM. ATCs were normalized to mean = 0 and SD = 1 beforehand.

Correlations, connectivity and functional relations

In addition to the ICs described above that had anatomical and temporal correspondences across subjects, there were numerous ICs that had the ‘looks’ of a functionally specialized region (e.g., a clean, often bilateral cluster of voxels in gray matter), that were, however, unique to only one subject, or whose ATCs across subjects showed no relation. These may be ICs that represent

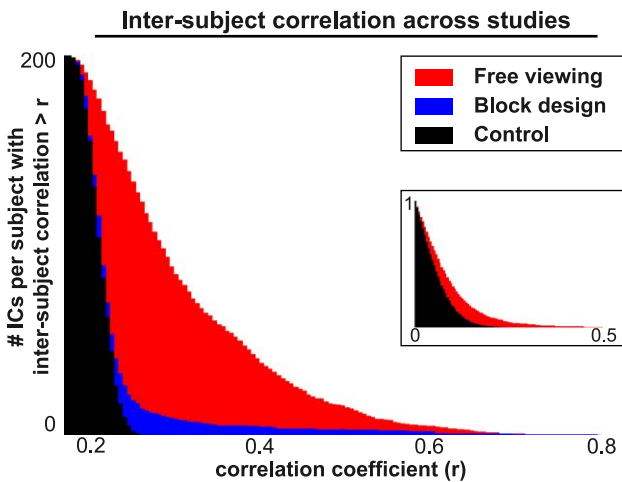


Fig. 9. More areas were differentially activated during free viewing than during conventional epoch stimulation. Shown are the cumulative distributions of maximal intersubject correlations among all IC-ATCs and the normalized cumulative distributions of intersubject correlations (inset) (see Methods). Both indicate how many ICs have ATCs that correlate across subjects and therefore provide a measure for the number of stimulus-driven regions that were differentially activated in each study. Free viewing data are shown in red and epoch data in blue. Shown in black is the expected distribution for simulated random ATCs as a baseline (white noise convolved with the HRF).

true, functionally specialized cortical subdivisions that were not stimulus-driven and therefore varied across subjects, for example, because they were involved in functions such as emotions, motor tasks, higher cognition or memory, which were not equally driven by the stimulus in every subject. Without any additional confirmation of their functional significance, for example, through spatial or temporal correlation across subjects, it cannot be excluded that these ICs represent artifacts of noninterest, such as those induced through movement, breathing, etc. For this reason, we did not discuss such ICs in further detail here, even if they are potentially false negatives. However, we briefly




Fig. 10. Hierarchical functional relationships between cortical areas revealed by the correlation of their ATCs. A clustering algorithm was used to group 15 representative ICs of one subject based on the correlation coefficients matrix of their ATCs ($d = 1 - |r|$). Branches were colored subsequently for graphical clarity. Ac = anterior cingulate, cer = cerebellum, FEF = frontal eye fields, m = motor cortex, pu = putamen.

propose a way to relate these ICs to a function, not by correlating them to the stimulus, but to other ICs within the same subject, which may have a related function, or which may have direct connections with each other and therefore related ATCs.

It is plausible to assume that cortical regions that have related functions, and also cortical regions that have direct anatomical connections (which may in most cases be the same areas), also have somewhat higher ATC correlations between each other during free-viewing conditions. Given the areas of an ICA-dissected human brain along with their activity time courses, it should therefore be possible to reveal the hierarchical organization of their proximity in terms of their function or their anatomical connections. Indeed, it has been shown that functionally connected cortical regions have specific and significant correlations between their ATCs during the ‘resting state’ of the human brain and that these correlations can be used to identify the connectivity of cortical areas in vivo (Arfanakis et al., 2000; Biswal et al., 1995; Cordes et al., 2000, 2001; Hampson et al., 2002). We hypothesize that correlations between cortical regions during free viewing can equally be used as a guide to functional or even anatomical connectivity between cortical regions, maybe even more so than ‘resting state’ correlations because free viewing leads to generally higher activity and to more specific activation than rest (Bartels and Zeki, 2004a,b). While we examine above concepts elsewhere, we address here briefly how correlations can be used to reveal a functional hierarchy between ICs. Correlations between ICs that can otherwise not be shown to be stimulus-driven can be used as an indicator that they may be functional rather than artifactual ICs. To illustrate the functional proximity of different ICs, we submitted ICs of a given subject to a clustering algorithm that grouped ICs in a hierarchical fashion using a distance matrix (d) based on the correlation (r) between their time courses: $d = (1 - |r|)$. This procedure therefore grouped areas according to their temporal similarity, which can be taken as an indicator for a similar function or for anatomical connections. We used a program originally designed to display the evolutionary phylogenetic relationship between different species with a distance matrix based on DNA sequences (PHYLIP package, with the KITSCH algorithm) to generate and draw the dendrograms (<http://evolution.genetics.washington.edu/phylip.html>); Felsenstein, 1981). Other clustering methods led to almost identical results. Fig. 10 shows how selected ICs of a single subject were temporally related. The ICs include some of above functional ICs and also some ICs that had no anatomically corresponding counterparts in other subjects. Although this is only an empirical demonstration, we note that similar results were obtained in all subjects, in that functionally related areas were grouped in the same subbranches of the resulting tree, and visual areas were segregated from other areas, as in the case shown. In this subject, we found the putamen to be highly correlated with the motor cortex; a subregion of the cerebellum to be grouped with the frontal eye fields, which are presumably both involved in eye movement control; and the cingulate gyrus to be coactive with a bilateral region in the lateral prefrontal cortex. While correlations among members of each group were relatively high, correlations between different groups were too low to attribute significance to the order in which they appeared in such trees. This method may therefore be useful for the identification of functional proximities between areas. Like a reverse correlation of activity to the stimulus, such hierarchical trees may be useful for the generation of hypotheses that may then be tested more formally.

Conclusion

We introduced here a new approach to mapping the functional subdivisions of the brain. It does not rely on a priori hypotheses or on histological markers, but establishes a new, area-specific marker, namely the characteristic activity time course (ATC) that is distinct between areas during exposure to natural conditions. Our analysis of a conventional epoch-design data set showed that the data-driven method of independent component analysis is capable of identifying regions that were activated differentially, which confirmed previous applications of ICA to fMRI (Bartels and Zeki, 2000a,b; Calhoun et al., 2001a, 2002; McKeown, 2000b; McKeown et al., 1998b; Zeki and Bartels, 1999). More importantly, it also showed that distinct areas exhibit differences in their ATCs even during exposure to identical stimuli, and that these characteristics were preserved across subjects and were area-specific. Our analysis of fMRI data collected during free-viewing conditions showed that distinct cortical areas exhibit distinct ATCs during free viewing, and that their temporal fingerprint is sufficiently specific to distinct areas that it can be used to map them separately. We further demonstrated that more functionally specialized areas are activated differentially during natural conditions than during conventional stimulation, and that distinct areas have a surprisingly high degree of temporal independence during natural conditions. We introduced methods that allow to prove the functional involvement of areas even if their function and the stimuli are unknown, by measuring their intersubject correlation and intersubject ranking, and we proposed ways to determine their function and their relation to other areas. Our findings are a further experimental step in confirming the principle of functional independence, which states that distinct functional subdivisions of the brain process features that vary independently, which is why they can be differentiated by their distinct ATCs, especially during exposure to more complex, natural conditions. Our dissections of different brains into their temporal subdivisions were consistent across subjects, suggesting that they revealed fundamental organizing principles of the human brain. We have therefore shown that temporal fingerprints of functionally distinct areas can be used to create a novel, time-based map of the brain, which we call *chronoarchitecture*. As it is a consequence of the functional and structural organization of the cortex, it does not reveal a cortical map that is distinct from that derived from anatomical tools or through classical functional mapping techniques. However, in contrast to the latter maps, it uses the marker of time, which is independent of human hypotheses about function, yet present and distinct in all cerebral subdivisions. Another advantage of this method is that the temporal relations between areas can provide further information about their functional and anatomical connectivity. Although we concentrated here on visual stimulation and visual areas, appropriate stimuli are likely to reveal further cortical and subcortical subdivisions in regions of the brain that are less well understood. Chronoarchitectonic maps may therefore guide the identification of unknown functional subdivisions in the human brain, which can then be further investigated using hypothesis-led experiments.

Acknowledgments

This work was supported by the Wellcome Trust, London. A.B. is supported by the Swiss National Science Foundation. We thank K. Friston, G. Rees, R. Perry and S. Shipp for their comments on

previous versions of the manuscript and A. Bork for making available fMRI data of the epoch study.

References

- Amari, S., 1999. Natural gradient learning for over- and under-complete bases. In: *ICA. Neural Comput.* 11 (8), 1875–1883.
- Arfanakis, K., Cordes, D., Haughton, V.M., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. *Magn. Reson. Imaging* 18 (8), 921–930.
- Bartels, A., Zeki, S., 1999. Can independent components analysis isolate cortical areas from fMRI data? *Soc. Neurosci. Abstr.* 766.5, 1930.
- Bartels, A., Zeki, S., 2000a. The architecture of the colour centre in the human visual brain: new results and a review. *Eur. J. Neurosci.* 12 (1), 172–193.
- Bartels, A., Zeki, S., 2000b. The neural basis of romantic love. *NeuroReport* 11 (17), 3829–3834.
- Bartels, A., Zeki, S., 2001. Chronoarchitecture: the dissection of the human brain into functional subdivisions by ICA analysis of fMRI data collected during free viewing of a movie. *Soc. Neurosci.* 620.15.
- Bartels, A., Zeki, S., 2003. Functional brain mapping during free viewing of natural scenes. *Hum. Brain Mapp.* 21 (2), 75–83 (online: 2003, print: 2004).
- Bartels, A., Zeki, S., 2004a. Brain dynamics during natural viewing conditions—a guide to cortical connectivity? (in preparation).
- Bartels, A., Zeki, S., 2004b. The chronoarchitecture of the human brain: functional anatomy based on natural brain dynamics and the principle of functional independence. In: Frackowiak, R., Friston, K., Frith, C., Dolan, R., Zeki, S. (Eds.), *Human Brain Function*, second ed. Academic Press, pp. 201–229.
- Beckers, G., Zeki, S., 1995. The consequences of inactivating areas V1 and V5 on visual-motion perception. *Brain* 118 (1), 49–60.
- Bell, A.J., Sejnowski, T.J., 1995. An information maximization approach to blind separation and blind deconvolution. *Neural Comput.* 7 (6), 1129–1159.
- Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34 (4), 537–541.
- Bork, A.C., Zeki, S., 1998. The cortical site for the generation of forms from motion. *NeuroImage* 7 (4 (2 of 3)), S329.
- Brodmann, K., 1909. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. J.A. Barth, Leipzig.
- Buchel, C., Holmes, A.P., Rees, G., Friston, K.J., 1998. Characterizing stimulus-response functions using nonlinear regressors in parametric fMRI experiments. *NeuroImage* 8 (2), 140–148.
- Calhoun, V.D., Adali, T., McGinty, V.B., Pekar, J.J., Watson, T.D., Pearlson, G.D., 2001a. fMRI activation in a visual-perception task: network of areas detected using the general linear model and independent components analysis. *NeuroImage* 14 (5), 1080–1088.
- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001b. A method for making group inferences from functional MRI data using independent component analysis. *Hum. Brain Mapp.* 14 (3), 140–151.
- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001c. Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Hum. Brain Mapp.* 13 (1), 43–53.
- Calhoun, V.D., Pekar, J.J., McGinty, V.B., Adali, T., Watson, T.D., Pearlson, G.D., 2002. Different activation dynamics in multiple neural systems during simulated driving. *Hum. Brain Mapp.* 16 (3), 158–167.
- Cavaglia, M., Dombrowski, S.M., Drazba, J., Vasanji, A., Bokesch, P.M., Janigro, D., 2001. Regional variation in brain capillary density and vascular response to ischemia. *Brain Res.* 910 (1–2), 81–93.
- Cordes, D., Haughton, V.M., Arfanakis, K., Wendt, G.J., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. *AJNR Am. J. Neuroradiol.* 21 (9), 1636–1644.
- Cordes, D., Haughton, V.M., Arfanakis, K., Carew, J.D., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2001. Frequencies contributing to functional connectivity in the cerebral cortex in “resting-state” data. *AJNR Am. J. Neuroradiol.* 22 (7), 1326–1333.
- d’Avossa, G., Shulman, G.L., Corbetta, M., 2003. Identification of cerebral networks by classification of the shape of BOLD responses. *J. Neurophysiol.* 90 (1), 360–371.
- DeYoe, E.A., Van Essen, D.C., 1985. Segregation of efferent connections and receptive field properties in visual area 2 of the macaque. *Nature* 317, 58–61.
- Duann, J.R., Jung, T.P., Kuo, W.J., Yeh, T.C., Makeig, S., Hsieh, J.C., Sejnowski, T.J., 2002. Single-trial variability in event-related BOLD signals. *NeuroImage* 15 (4), 823–835.
- Felsenstein, J., 1981. Evolutionary trees from DNA sequences: a maximum likelihood approach. *J. Mol. Evol.* 17 (6), 368–376.
- Ffytche, D.H., Howard, R.J., Brammer, M.J., David, A., Woodruff, P., Williams, S., 1998. The anatomy of conscious vision: an fMRI study of visual hallucinations. *Nat. Neurosci.* 1 (8), 738–742.
- Flechsig, P., 1901. Developmental (myelogenetic) localisation of the cerebral cortex in the human subject. *Lancet* 2, 1027–1029.
- Frackowiak, R., Friston, K., Frith, C., Dolan, R., Mazziotta, J.C. (Eds.), 1997. *Human Brain Function*. Academic Press.
- Friston, K.J., 1998. Modes or models: a critique on independent component analysis for fMRI. *Trends Cogn. Sci.* 2, 373–374.
- Friston, K.J., Frith, C.D., Frackowiak, R.S.J., Turner, R., 1995a. Characterizing dynamic brain responses with fMRI—a multivariate approach. *NeuroImage* 2 (2 Pt 1), 166–172.
- Friston, K.J., Holmes, A.P., Poline, J.B., Grasby, P.J., Williams, S.C.R., Frackowiak, R.S.J., Turner, R., 1995b. Analysis of fMRI time-series revisited. *NeuroImage* 2 (1), 45–53.
- Gu, H., Engelen, W., Feng, H., Silbersweig, D.A., Stern, E., Yang, Y., 2001. Mapping transient, randomly occurring neuropsychological events using independent component analysis. *NeuroImage* 14 (6), 1432–1443.
- Hampson, M., Peterson, B.S., Skudlarski, P., Gatenby, J.C., Gore, J.C., 2002. Detection of functional connectivity using temporal correlations in MR images. *Hum. Brain Mapp.* 15 (4), 247–262.
- Harrison, R.V., Harel, N., Panesar, J., Mount, R.J., 2002. Blood capillary distribution correlates with hemodynamic-based functional imaging in cerebral cortex. *Cereb. Cortex* 12 (3), 225–233.
- Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., Malach, R., 2003. Brain show: functional organization of the human cortex under free viewing of natural vision. *Soc. Neurosci [Abstract 659.2]*.
- Henson, R.N., Price, C.J., Rugg, M.D., Turner, R., Friston, K.J., 2002. Detecting latency differences in event-related BOLD responses: application to words versus nonwords and initial versus repeated face presentations. *NeuroImage* 15 (1), 83–97.
- Hubel, D.H., Livingstone, M.S., 1987. Segregation of form, color and stereopsis in primate area 18. *J. Neurosci.* 7, 3378–3415.
- Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* 17 (11), 4302–4311.
- Kleinschmidt, A., Büchel, C., Zeki, S., Frackowiak, R.S.J., 1998. Human brain activity during spontaneously reversing perception of ambiguous figures. *Proc. R. Soc. London, Ser. B*, vol. 265, pp. 2422–2427.
- Livingstone, M.S., Hubel, D.H., 1984. Anatomy and physiology of a color system in the primate visual cortex. *J. Neurosci.* 4, 309–356.
- Lumer, E.D., Friston, K.J., Rees, G., 1998. Neural correlates of perceptual rivalry in the human brain. *Science* 280 (5371), 1930–1934.
- Makeig, S., Jung, T.P., Bell, A.J., Ghahremani, D., Sejnowski, T.J., 1997a. Blind separation of auditory event-related brain responses into independent components. *Proc. Natl. Acad. Sci. U. S. A.* 94 (20), 10979–10984.
- Makeig, S., Jung, T.P., Bell, A.J., Ghahremani, D., Sejnowski, T.J.,

- 1997b. Blind separation of auditory event-related brain responses into independent components. *Proc. Natl. Acad. Sci. U. S. A.* 94 (20), 10979–10984.
- Makeig, S., Westerfield, M., Jung, T.P., Enghoff, S., Townsend, J., Courchesne, E., Sejnowski, T.J., 2002. Dynamic brain sources of visual evoked responses. *Science* 295 (5555), 690–694.
- Malach, R., Reppas, J.B., Benson, R.R., Kwong, K.K., Jiang, H., Kennedy, W.A., Ledden, P.J., Brady, T.J., Rosen, B.R., Tootell, R.B.H., 1995. Object-related activity revealed by functional magnetic-resonance-imaging in human occipital cortex. *Proc. Natl. Acad. Sci. U. S. A.* 92 (18), 8135–8139.
- McKeown, M.J., 2000a. Detection of consistently task-related activations in fMRI data with hybrid independent component analysis. *NeuroImage* 11, 24–35.
- McKeown, M.J., 2000b. Detection of consistently task-related activations in fMRI data with hybrid independent component analysis. *NeuroImage* 11 (1), 24–35.
- McKeown, M.J., Sejnowski, T.J., 1998. Independent component analysis of fMRI data: examining the assumptions. *Hum. Brain Mapp.* 6, 368–372.
- McKeown, M.J., Jung, T.P., Makeig, S., Brown, G., Kindermann, S.S., Lee, T.W., Sejnowski, T.S., 1998a. Spatially independent activity patterns in functional MRI data during the Stroop color naming task. *Proc. Natl. Acad. Sci. U. S. A.* 95 (3), 803–810.
- McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J., 1998b. Analysis of fMRI data by blind separation into independent spatial components. *Hum. Brain Mapp.* 6 (3), 160–188.
- McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.S., 1998c. Analysis of fMRI data by blind separation into independent spatial components. *Hum. Brain Mapp.* 6 (3), 160–188.
- Moutoussis, K., Zeki, S., 1997. A direct demonstration of perceptual asynchrony in vision. *Proc. R. Soc. London, Ser. B* 264 (1380), 393–399.
- Nakada, T., Suzuki, K., Fujii, Y., Matsuzawa, H., Kwee, I.L., 2000. Independent component-cross correlation-sequential epoch (ICS) analysis of high field fMRI time series: direct visualization of dual representation of the primary motor cortex in human. *Neurosci. Res.* 37, 237–244.
- Nybakken, G.E., Quigley, M.A., Moritz, C.H., Cordes, D., Haughton, V.M., Meyerand, M.E., 2002. Test–retest precision of functional magnetic resonance imaging processed with independent component analysis. *Neuroradiology* 44 (5), 403–406.
- Schiller, P.H., 1996. On the specificity of neurons and visual areas. *Behav. Brain Res.* 76 (1–2), 21–35.
- Schiller, P.H., 1997. Past and present ideas about how the visual scene is analyzed by the brain. In: Rockland, K.S., Kaas, J.H., Peters, A. (Eds.), *Extrastriate Cortex in Primates*. Plenum, New York, pp. 59–90.
- Schmolecky, M.T., Wang, Y., Hanes, D.P., Thompson, K.G., Leutgeb, S., Schall, J.D., Leventhal, A.G., 1998. Signal timing across the macaque visual system. *J. Neurophysiol.* 79 (6), 3272–3278.
- Shipp, S., Zeki, S., 1985. Segregation of pathways leading from area V2 to areas V4 and V5 of macaque monkey visual cortex. *Nature* 315, 322–325.
- Stone, J.V., Porrill, J., Porter, N.R., Wilkinson, I.D., 2002. Spatiotemporal independent component analysis of event-related fMRI data using skewed probability density functions. *NeuroImage* 15 (2), 407–421.
- Suzuki, K., Kiryu, T., Nakada, T., 2002. Fast and precise independent component analysis for high field fMRI time series tailored using prior information on spatiotemporal structure. *Hum. Brain Mapp.* 15 (1), 54–66.
- Talairach, J., Tournoux, P., 1988. *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme, Stuttgart.
- Vogt, C., Vogt, O., 1919. Allgemeinere Ergebnisse unserer Hirnforschung. Vierte Mitteilung. Die physiologische Bedeutung der architektonischen Rindenfelderung auf Grund neuer Rindenreizungen. *J. Psychol. Neurol.* 25, 399–462.
- Zeki, S., 1978. Functional specialization in the visual cortex of the monkey. *Nature* 274, 423–428.
- Zeki, S., Bartels, A., 1998. The asynchrony of consciousness. *Proc. R. Soc. London, Ser. B* 265, 1583–1585.
- Zeki, S., Bartels, A., 1999. The clinical and functional measurement of cortical (in-) activity in the visual brain, with special reference to the two subdivisions (V4 and V4a) of the human colour centre. *Philos. Trans. R. Soc. London, Ser. B* 354, 1371–1382.
- Zheng, D., LaMantia, A.S., Purves, D., 1991. Specialized vascularization of the primate visual cortex. *J. Neurosci.* 11 (8), 2622–2629.