Analysis of fMRI Data by Blind Separation Into Independent Spatial Components

Martin J. McKeown,[1*] **Scott Makeig,**[2,3] **Greg G. Brown,**[5] **Tzyy-Ping Jung,**[1]
Sandra S. Kindermann,[5] **Anthony J. Bell,**[1] **and Terrence J. Sejnowski**[1,4]

[1]*Howard Hughes Medical Institute, Computational Neurobiology Laboratory,*
Salk Institute for Biological Studies, La Jolla, California 92186-5800
[2]*Cognitive Psychophysiology Laboratory, Naval Health Research Center, San Diego,*
California 92186-5122
[3]*Department of Neurosciences, School of Medicine, University of California at San Diego,*
La Jolla, California 92093
[4]*Department of Biology, University of California at San Diego, La Jolla, California 92093*
[5]*Department of Psychiatry, School of Medicine, University of California at San Diego,*
La Jolla, California 92093

◆ ──────────────────────────────── ◆

Abstract: Current analytical techniques applied to functional magnetic resonance imaging (fMRI) data require a priori knowledge or specific assumptions about the time courses of processes contributing to the measured signals. Here we describe a new method for analyzing fMRI data based on the independent component analysis (ICA) algorithm of Bell and Sejnowski ([1995]: Neural Comput 7:1129–1159). We decomposed eight fMRI data sets from 4 normal subjects performing Stroop color-naming, the Brown and Peterson word/number task, and control tasks into spatially independent components. Each component consisted of voxel values at fixed three-dimensional locations (a component ''map''), and a unique associated time course of activation. Given data from 144 time points collected during a 6-min trial, ICA extracted an equal number of spatially independent components. In all eight trials, ICA derived one and only one component with a time course closely matching the time course of 40-sec alternations between experimental and control tasks. The regions of maximum activity in these consistently task-related components generally overlapped active regions detected by standard correlational analysis, but included frontal regions not detected by correlation. Time courses of other ICA components were transiently task-related, quasiperiodic, or slowly varying. By utilizing higher-order statistics to enforce successively stricter criteria for spatial independence between component maps, both the ICA algorithm and a related fourth-order decomposition technique (Comon [1994]: Signal Processing 36:11–20) were superior to principal component analysis (PCA) in determining the spatial and temporal extent of task-related activation. For each subject, the time courses and active regions of the task-related ICA components were consistent across trials and were robust to the addition of simulated noise. Simulated movement artifact and simulated task-related activations added to actual fMRI data were clearly separated by the algorithm. ICA can be used to distinguish between nontask-related signal components, movements, and other artifacts, as well as consistently or transiently task-related fMRI activations, based on only weak

*Correspondence to: Dr. M.J. McKeown, Computational Neurobiology Laboratory, Salk Institute for Biological Studies, 10010 North Torrey Pines Road, La Jolla, CA 92037-1099. E-mail: martin@salk.edu

assumptions about their spatial distributions and without a priori assumptions about their time courses. ICA appears to be a highly promising method for the analysis of fMRI data from normal and clinical populations, especially for uncovering unpredictable transient patterns of brain activity associated with performance of psychomotor tasks. *Hum. Brain Mapping 6:160–188, 1998.* © 1998 Wiley-Liss, Inc.

Key words: functional magnetic resonance imaging; independent component analysis, higher-order statistics

◆ ═══ ◆

INTRODUCTION

Many current functional magnetic resonance imaging (fMRI) experiments use a block design in which the subject is instructed to perform experimental (E) and control (C) tasks in an alternating sequence of 20–40-sec blocks (e.g., CECECEC...). During such a trial, signals from thousands of volume elements (voxels) in each of several brain slices are typically acquired every 1–3 sec. The resultant time series recorded for each voxel may contain a complicated mixture of high- and low-frequency activity (Fig. 1), which is most probably produced by a medley of local or spatially distributed processes, including task-related and nontask-related hemodynamic brain tissue activations as well as motion or machine artifacts. This tangled mixture of signals presents a formidable challenge for analytical methods attempting to tease apart task-related changes from the disparate time courses of 5,000–25,000 voxels.

Changes in fMRI signal (including blood oxygen level-dependent (BOLD) contrast [Ogawa et al., 1992]) related to alternating performance of experimental and control tasks have been analyzed by a number of techniques, including subtraction, correlation, and time-frequency analyses, and have been tested statistically using t-tests [Kwong et al., 1992], analysis of variance/covariance (ANOVA/ANCOVA) (Friston, 1996), and nonparametric Komolgorov-Smirnov tests [Stuart and Ord, 1991; Kwong, 1995].

Subtraction or, more generally, correlation techniques [Bandettini et al., 1993] are based on the assumption that voxels indexing brain regions participating in the cognitive processing of the given experimental and control tasks should show different fMRI signal levels during the performance of these tasks. Correlation techniques exploit a priori knowledge of the expected time course of task-related changes in the signal to determine their intensity and spatial extent. A *reference function* is created by convolving the block design of the behavioral experiment (CECECEC...) with a fixed model of the hemodynamic response function (an estimate of the fMRI signal changes evoked by a brief burst of neural activity). This reference function is then correlated with the time series recorded from each voxel. Those voxels, whose signals are positively correlated with the reference function above a preselected threshold, are designated "areas of activation." Although this method is both computationally simple and reasonably effective, it has several major drawbacks. Even in areas of activation, the task-related signal changes are typically small (<10%), suggesting that other time-varying phenomena must produce the bulk of the measured signals. These phenomena can be conceptualized as multiple concurrent "*component processes,*" each having a separate time course and spatial extent and each producing simultaneous changes in the fMRI signals of many voxels. Other component processes may not be completely uncorrelated with task-related changes, and so may tend to mask the effects of activations related to task-performance, reducing the sensitivity and specificity of correlational analysis. If the nontask-relevant component processes are monotonic and linear, simple linear detrending [Bandettini et al., 1993] can be expected to enhance the accuracy of correlational analysis. However, the time courses of processes related to changes in arousal, task strategy, head position, machine artifacts, or other endogenous processes occurring during a trial may not resemble simple linear or nonlinear functions.

More general ANOVA-like approaches [Friston, 1996], including statistical parametric mapping (SPM) [Friston, 1995], test the signal at each voxel using univariate measures (e.g., t-tests, or f-tests) under the null hypothesis that the values are distributed under a known probability distribution (typically Gaussian). Voxels in which the signal difference between the task and control conditions exceeds a predefined level of significance are selected as active, resulting in a distributed spatial image giving anatomical areas of significant task-related activation difference. Using this technique, it is possible to test multiple factors that may contribute to changes in the fMRI signals in addition to the task design. However, ANOVA-like methods are based on the assumptions, tenuous for fMRI data, that: 1) the observations have a known distribution (e.g., Gaussian), 2) the variances and covariances between repeated measurements are equal, 3) the time courses
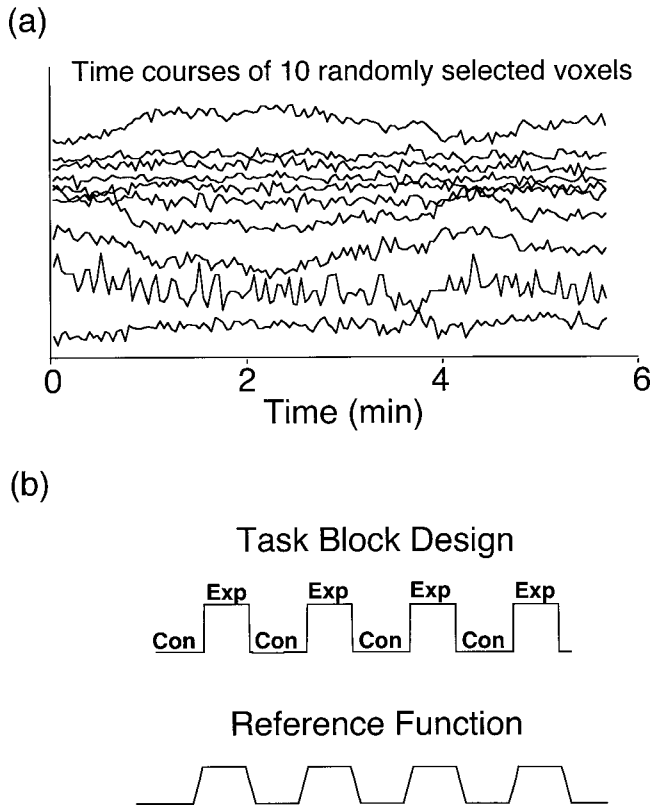
(a)

Time courses of 10 randomly selected voxels

(b)

Task Block Design

Reference Function

Figure 1.
BOLD signal complexity and task reference function. **a:** Time courses of 10 randomly selected voxels from a 6-min fMRI trial of the Stroop color-naming task illustrate the typical complexity of BOLD signals. **b:** Convolving an a priori estimate of the hemodynamic response function with the square-wave function representing the task block structure of the trial, alternating experimental (Exp) and control (Con) blocks (upper trace) produce the reference function for the trial (bottom trace).

nent sources, which may not be valid, although wavelet techniques currently being explored might possibly relax this requirement [Brammer et al., 1997].

Correlational, time-frequency, and ANOVA-based methods share an inherent weakness common to univariate techniques currently used for analysis of fMRI data: they do not attempt to extract the *intrinsic* structure of the data. This could be a particularly significant drawback in cases where accurate a priori models of fMRI signal changes in response to experimental events are not known or may not be constant across all voxels, e.g., in data from patient populations with pathological brain conditions, or from subjects performing complex learning tasks. Another drawback of ANOVA-based and correlational measures is that they typically require grouping or averaging data over several task/control blocks. This reduces their sensitivity for detecting transient task-related changes in the fMRI signal, and makes them insensitive to significant changes not consistently time-locked to the task block design. These could include changes in strategy by the subject during the test period, changes associated with learning or habituation of task performance, with fatigue, or with other processes whose time courses cannot be predicted in advance by the experimenter. Univariate techniques also ignore relationships between voxels, hindering the detection of brain regions acting as functional units during the experiment.

Principal component analysis (PCA) has been proposed as a way to isolate functional patterns in functional imaging data [Moeller et al., 1991]. This technique first measures the tendency of signals at all possible pairs of voxels to covary, and then finds the orthogonal spatial patterns or *eigenimages* capturing the greatest variance in the data. The first eigenimage represents the largest source of variance between pairs of voxels, the second eigenimage represents the largest source of residual variance orthogonal to the first eigenimage, and so on. Normally, the number of principal components required to adequately represent the data to a specified level of accuracy is much smaller than the original dimension of the data [Jackson, 1991], and thus PCA can provide a useful method for reducing data dimensionality. However, if task-related fMRI changes are only a small part of the total signal variance, retaining the orthogonal eigenimages capturing the greatest variance in the data may reveal little information about task-related activations or other processes of interest. Additionally, if during an fMRI experiment numerous voxels become simultaneously activated, component analysis methods based solely on voxel-pair relationships or covariances may

of different factors affecting the variance of the fMRI signal can be reliably estimated in advance, and 4) the signals at different voxels are independent. Signal distributions can be made more Gaussian by spatial and temporal smoothing, but this smoothing also degrades the temporal and spatial resolution of the data.

Time-frequency analyses describe the signal recorded from each voxel in the frequency domain and may be useful for distinguishing between physiological pulsatile and other repetitive artifacts known to be present in fMRI data [Mitra et al., 1997]. Such techniques assume that signal change produced by task performance and other sources of physiological interest have frequency spectra different from other causes of variability in the data. Many of these techniques assume periodicity in the time courses of the compo-

not capture their overall patterns of association. These shortcomings suggest the desirability of a general fMRI analytical technique capable of extracting the intrinsic spatiotemporal structure of the data without the aforementioned limitations associated with PCA and other existing analytical tools.

Here we describe a new technique for the analysis of fMRI data based on the statistical method of independent component analysis (ICA) [Comon, 1994; Bell and Sejnowski, 1995]. It potentially allows the extraction of both transient and consistently task-related, as well as physiologically-relevant nontask-related, and various artifactual components of the observed fMRI signals.

INDEPENDENT COMPONENT ANALYSIS

Functional organization of the brain is based on two complementary principles, *localization* and *connectionism* [Phillips et al., 1984]. *Localization* implies that each psychomotor function is performed principally in a small set of brain areas. This principle derives originally from clinical experience where a restricted locus of damage to the nervous system could usually be inferred from a specific pattern of deficits demonstrated by a subject [Gardner, 1975]. Occasionally, the locus of the lesion cannot accurately be directly determined by the pattern of deficits, as in the clinical "disconnection syndromes" (e.g., alexia without agraphia [Duffield et al., 1994; Quint and Gilmore, 1992] and pure word deafness [Takahashi et al., 1992]), because the lesion interrupts connections between macroscopic loci required to perform some psychomotor task. This demonstrates the complementary principle of *connectionism* that posits that the brain regions involved in a given psychomotor function may be widely distributed, and thus the brain activity required to perform a given task may be the functional integration of activity in multiple macroscopic loci or distinct brain systems (this is a different sense of the term "connectionism" from that used to describe neural network models).

Consistent with these principles, we suggest that the multifocal brain areas activated by performance of a psychomotor task should be unrelated to the brain areas whose signals are affected by artifacts, such as physiological pulsations, subtle head movements, and machine noise which may dominate fMRI experiments. Each of these separate processes may be represented by one or more *spatially-independent* components, each associated with a single time course of enhancement and/or suppression and a component map (Fig. 2). We assume the component maps, each specified by a spatial distribution of fixed values (one


Figure 2.
Schematic of fMRI data decomposed into independent components. Each independent component produced by the ICA algorithm consists of a spatial distribution of voxel values ("component map"), and an associated time course of activation. The four schematic component maps show voxels participating most actively in each of four hypothetical components. Under ICA, the signal observed at a given voxel is modeled as a sum of the contributions of all the independent components. The amount each component contributes to the data is determined by the outer product of the voxel values in its component map with the activation values in its time course. Note that active areas of statistically independent map value distributions may be partially overlapping.

at each voxel), represent possibly overlapping, multifocal brain areas of statistically dependent fMRI signal influence. Furthermore, we presume that the component map distributions are *spatially independent,* and hence uniquely specified. This means that if $p_k(C_k)$ specifies the probability distribution of the voxel values C_k in the k^{th} component map, then the joint probability distribution of all n components factorizes:

$$p(C_1, C_2, \ldots, C_n) = \prod_{k=1}^{n} p_k(C_k) \qquad (1)$$

where each of the component maps C_k is a vector (C_{ki}, i = 1, 2, ... M), and M is the number of voxels.

Note that this is a much stronger criterion than saying that the voxel values between pairs of components are merely uncorrelated, i.e.,

$$C_i \cdot C_j = \sum_{k=1}^{M} C_{ik}C_{jk} = 0, \quad \text{for } i \neq j \quad (2)$$

since Equation (1) implies that higher-order correlations are also zero.

The maps will be independent if active voxels in the maps are sparse and mostly nonoverlapping [McKeown et al., 1998], although in general some overlap will occur. We further assume that the observed fMRI signals are the linear sum of the contributions of the individual component processes at each voxel. With these assumptions, the fMRI signals recorded during the performance of psychomotor tasks can be decomposed into a number of independent component maps and their associated component activation waveforms, using the ICA algorithm given below (Figs. 2, 3). No a priori assumptions need be made about the time courses of activation of the different components, or whether a given component is activated by specific psychophysiological systems or is related to machine noise or other artifacts.

These ideas can be expressed rigorously by writing a matrix equation relating the component maps and their time courses to the measured fMRI signals. If the map voxel values for each of the components are known and placed in separate rows of matrix C_{ki}, then a mixing matrix, M_{jk}, can specify the time-varying contributions of each component map to the measured fMRI signals (Fig. 3):

$$X_{ji} = \sum_{k=1}^{N} M_{jk}C_{ki}. \quad (3)$$

Decomposing observed fMRI signals into statistically independent component maps without prior knowledge of their spatial extents or time courses of activation is a "blind separation" problem [Jutten and Herault, 1991]. The independent component analysis (ICA) algorithm of Bell and Sejnowski [1995], an iterative unsupervised neural network learning algorithm based on information-theoretic principles, can perform blind separation of input data into the linear sum of time-varying modulations of maximally independent component maps. The ICA algorithm iteratively determines the unknown *unmixing* matrix W, a possibly linearly scaled and permuted version of the *inverse* of the *mixing* matrix, M, from which the compo-

Figure 3.
fMRI data as a mixture of independent components. The mixing matrix M specifies the relative contribution of each component at each time point. ICA finds an *unmixing* matrix that separates the observed component mixtures into the independent component maps and time courses.

nent maps and time courses of activation can be computed.

The matrix of component maps is computed by multiplying the observed data matrix X by W,

$$C_{ij} = \sum_{k=1}^{N} W_{ik}X_{kj} \quad (4a)$$

where W_{ik} is the unmixing matrix derived from ICA, C_{ij} is the value of the j^{th} voxel in the i^{th} component map, X_{kj} is the k^{th} time point of the j^{th} voxel, and the summation runs over the N time points of the fMRI input data. In matrix notation, this can be written simply as

$$C = WX \quad (4b)$$

where X is the fMRI signal data matrix, W is the unmixing matrix, and C is the matrix of component map voxel values. Note that W is a square matrix of full rank, so its inverse W^{-1} is well-defined. Although a nonlinearity is used by the algorithm in the *determination* of W (see below), W itself provides a *linear* decomposition of the data.

Reconstruction of the data from the independent components is accomplished by

$$X'_{ij} = \sum_{k=1}^{N} W^{-1}_{ik} C_{kj} \qquad (5a)$$

where X'_{ij} is the reconstructed data at the i^{th} time point of the j^{th} voxel, and the summation runs over the N time points of the fMRI input data. In matrix notation,

$$X' = W^{-1}C. \qquad (5b)$$

The data can be perfectly reconstructed when $W^{-1} = M$, i.e., $X' = W^{-1}C = MC = X$. The first column of W^{-1} gives the time course of modulation of the first component map, the second column gives the time course of the second component map, and so on. The ICA method can extract a number of independent components up to the number of time points in the data, each having a map that does not change during the course of a trial and a unique associated time course of activation. The distributions of voxel values in the component maps, C, are as statistically independent as possible, while the component time courses (contained in W^{-1}) may be correlated. The order of the rows of W, and hence of the calculated ICA components, is not meaningful and may vary between repeated analyses of the same data. It is useful therefore to rank order the components by the extent of their contribution to the original data.

Rank ordering of the components is complicated by the fact that the different ICA component time courses contained in W^{-1} are, in general, nonorthogonal so that, unlike PCA, the variances explained by each component will not sum to the variance of the original data. The *contribution* each component makes to the magnitude of the original data, γ_i, can be estimated by the root mean square (RMS) of the data set reconstructed solely from this component, i.e., from Equation (5b) with C having one nonzero row corresponding to the appropriate component. Alternatively, the contribution can be considered the RMS error introduced per data point when the data is reconstructed without this component. Thus:

$$\gamma_i = \frac{1}{NM} \left(\sum_{j=1}^{N} \sum_{k=1}^{M} A^2_{jk} \right)^{\frac{1}{2}} \qquad (6)$$

where γ_i is the contribution to the data from the i^{th} component, N is the number of time points, M is the number of brain voxels, and A^i_{jk} is an (N by M) matrix computed from the outer product of the i^{th} component map and i^{th} column of W^{-1}, i.e.,

$$A^i_{jk} = W^{-1}_{ji} C_{ik}. \qquad (7)$$

Each ICA component map is described by a distribution of values, one for each voxel. These values represent the relative amount a given voxel is modulated by the activation of that component. To find and display voxels contributing significantly to a particular component map, the map values may be scaled to z-scores (the number of standard deviations from the map mean). Voxels whose absolute z-scores are greater than some threshold (e.g., $|z| > 2$) can be considered to be "*active*" voxels for that component. In this case, the z-scores are used for descriptive purposes and have no definite statistical interpretation. Negative z-scores indicate voxels whose fMRI signals are modulated *opposite* to the time course of activation for that component.

In summary, unlike other methods of fMRI analysis that begin with a matrix of Pearson product-moment correlations [Moeller et al., 1991], ICA utilizes a much stronger criterion for statistical independence. ICA decomposes the observed fMRI data into maps of activities that are as spatially independent as possible and provides a unique representation of the data (up to scaling and permutation).

THE ICA ALGORITHM

Under the assumption that component processes can be represented by differentially activated spatially sparse and spatially independent maps, and that the sum of their activations equals the observed data, an unmixing matrix W can be determined using a statistical method based on the "infomax" principle [Bell and Sejnowski, 1995]. In information theory, the probability of a message and its informational content are inversely related. More formally, the mean uncertainty or *entropy* associated with a set of messages in discrete form is

$$H(X) = -\sum_k p_k \log p_k \qquad (8)$$

where p_k is the probability of the k_{th} event.

The joint entropy of two variables is defined by $H(X, Y) = H(X) + H(Y) - I(X, Y)$, where $I(X, Y) = H(X) - H(X|Y)$ is the *mutual information,* interpreted as the

redundancy between X and Y or, alternatively, as the reduction in uncertainty of one variable (e.g., X) due to the observation of the other variable (Y). ICA attempts to maximize the joint entropy of suitably transformed (see below) component maps, and in so doing reduces the redundancy between the distributions of map values for different components. This, in effect, results in blind separation of the recorded fMRI signals into spatially independent components. See Bell and Sejnowski [1995] for a more detailed discussion of the training process.

In brief, the algorithm initializes W to the identity matrix (I), then iteratively attempts to maximize H(y), where y = g(C), C = WX$_s$, and g() is a specified nonlinear function. Here, X$_s$ is a "sphered" version of the data matrix defined by

$$X_s = Px, \tag{9}$$

where

$$P = 2\langle XX^T \rangle^{-1/2} \tag{10}$$

and $\langle XX^T \rangle$ is the N × N covariance matrix of the data matrix X.

The nonlinear function g(), which provides necessary higher-order statistical information, is chosen here to be the logistic function

$$g(C_i) = \frac{1}{1 + e^{-C_i}} \tag{11}$$

which biases the algorithm towards finding spatially sparse component maps with relatively few highly active voxels [McKeown et al., 1998].

The elements of W are updated using small batches of data vectors drawn randomly from [X$_s$] without substitution, according to:

$$\Delta W = -\epsilon \left(\frac{\partial H(y)}{\partial W} \right) W^T W = \epsilon (I + \hat{y}C^T)W \tag{12}$$

where ϵ is a learning rate (typically near 0.01) and the vector $\hat{y}$ has elements

$$\hat{y}_i = \frac{\partial}{\partial C_i} \ln \left(\frac{\partial y_i}{\partial C_i} \right) = (1 - 2y_i). \tag{13}$$

The W^{T}W term in Equation (12), first proposed by Amari et al. [1996], avoids matrix inversions and speeds convergence. During training, the learning rate is reduced gradually until the weight matrix W stops changing appreciably (e.g., root mean square change for all elements <10^{-6}).

COMON'S FOURTH-ORDER TECHNIQUE FOR INDEPENDENT COMPONENT ANALYSIS

Common [1994] defined the concept of independent component analysis (ICA) as determining a linear transformation, W, such that application to decorrelated input data resulted in approximately maximal statistical independence between outputs (in this case, component maps). He demonstrated that W could be estimated by maximizing a contrast function, $\phi(W)$, using a computationally intensive method that iteratively updated W using all data points to estimate $\phi(W)$ at each iteration.

This method finds a linear transformation, W, which maximizes

$$\phi(W) = \sum_{j=1}^{N} (k_4^j)^2 \tag{14}$$

where

$$k_4^j = m_4^j - 4m_3^j m_1^j - 3(m_2^j)^2 + 12m_2^j(m_1^j)^2 - 6(m_1^j)^4. \tag{15}$$

Here, m_i is the ith moment defined as

$$m_i^j = \int_{-\infty}^{\infty} C_j^i p(C_j) dC_j. \tag{16}$$

C_j is the jth component map computed from C = WX, and X is the N by V data matrix.

ILLUSTRATIVE THOUGHT EXPERIMENT

To assist in the reader's appreciation of the conceptual differences between ICA, PCA, and correlation analyses, we propose and analyze a two-dimensional "thought experiment" (Fig. 4). Imagine an fMRI data set that is the sum of the contributions of just two spatially-independent component processes (IC1 and IC2). The data are recorded at two separate time points during an experimental session. At time point t = 1, the subject is performing an experimental task, while time point t = 2 occurs during a control task condition. The two component processes, portrayed schematically in Figure 4a, are primarily active during the control and experimental task periods, respectively. Component IC2 is mostly task-related, since it is highly active at t = 1 and only weakly active at t = 2.

An fMRI "Thought Experiment"

(a)

(b)
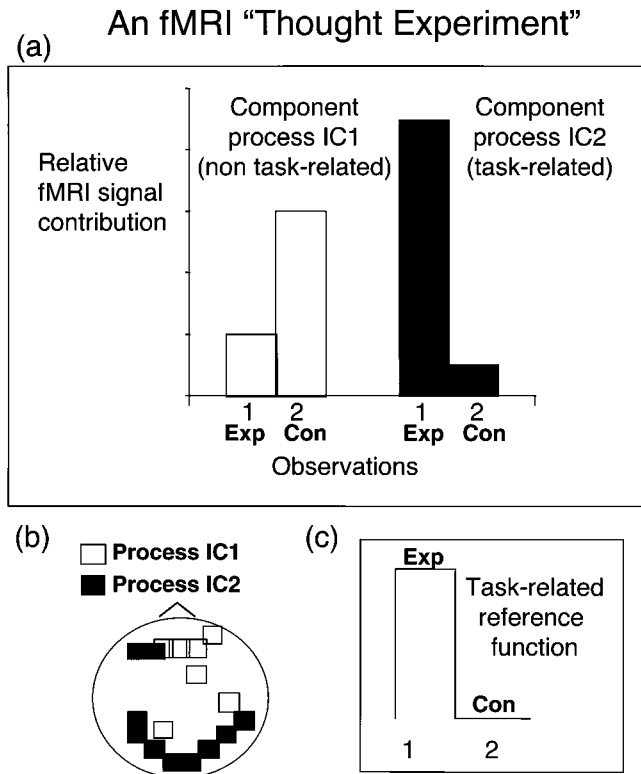☐ **Process IC1**
■ **Process IC2**

(c)

Figure 4.
Simulated experiment. **a:** A simple "thought experiment" to demonstrate differences between ICA, PCA, and correlation analysis methods. A hypothetical fMRI data set is the sum of the activity of just two spatially-independent processes (IC1 and IC2) recorded at two observation times (t = 1, experimental; t = 2, control). We assume that process IC2 is mostly task-related (e.g., it is highly active at t = 1 and only weakly active at t = 2), and process IC1 (e.g., representing endogenous activity or machine artifact) is more active at t = 2. We further assume that the distributions of voxel values of the two components specifying the locations of the processes are independent of one another. **b:** The voxels with the largest map values in the two hypothetical component distributions are *active* voxels of the components. **c:** The simplest reference function useful for detecting task-related activations using correlation analysis is active (=1) during the experimental task and silent (=0) during the control task.

Component IC1 (representing either endogenous activity or machine artifact) is somewhat more active at t = 1 than at t = 2. We assume that the distributions of voxel values for the two components are independent of one another, with fairly small and discrete sets of active voxels (such as those indicated for the cartoon head in Fig. 4b). Here, a simple reference function for detecting task-related brain areas via correlation (Fig. 4c) will have the values 1 (= "on") at t = 1 and 0 (= "off") at t = 2.

Figure 5 (top) shows a scatter plot of the hypothetical fMRI data. Here, for each voxel the signal value recorded at time t = 1 is plotted against its value at t = 2. The relative activations of components IC1 and IC2 (Fig. 4) appear in Figure 5 as fixed vector directions. The assumption that component processes IC1 and IC2 are spatially independent implies that the data points will tend to fill along each of the vectors labeled IC1 and IC2. Note that the distribution of data values at t = 1 is correlated with the data distribution at t = 2. Thus, the marginal probability distributions of the data in its current form are not uniform, and the data distribution cannot have maximum entropy.

Applying a suitable linear transformation, W, to the data transforms it into the rectangular data distribution shown in Figure 5b. Under W, the IC1 and IC2 vectors in the upper plot are mapped into the orthogonal basis axes IC1′ and IC2′, which "unmix" the contributions of processes IC1 and IC2 to the data. The assumed spatial independence of IC1 and IC2 means that the transformed data are then rectangular, and thus have higher entropy than the original data. If each component map is sparse, with relatively few large values, passing the transformed data through a sigmoidal nonlinearity g() (Fig. 5c) will more evenly spread out the data within the rectangle, producing a data distribution g(WX) that has still larger entropy. The ICA algorithm attempts to find directions IC1 and IC2 by iteratively adjusting W so as to maximize the entropy of the resulting transformed distribution g(WX) (Fig. 5c). Note also that the linear transform, W, is in general unique only up to scaling and permutation (e.g., W, might switch the orders of IC1 and IC2).

Active voxels in the IC1 and IC2 component maps (e.g., voxels that would be indicated in a head image such as Fig. 4b) are those that project most strongly on vectors IC1′ and IC2′ (e.g., the voxels inside the solid parallelograms in Fig. 5a). Active voxels according to correlation of the data with the assumed reference function are those whose projections onto the reference function direction (COR) exceed some threshold (e.g., those inside the dashed rectangle in Fig. 5a). Unlike ICA, principal component analysis (PCA) finds orthogonal directions of *maximum variance* in the data. The eigenvector associated with the first principal component points in the direction of maximum variance of the data (PC1 in Fig. 5a). In general, this has no specific relationship to the directions (i.e., time courses) of the independent components. Active voxels in the PC1 direction are, e.g., those inside the tilted dotted rectangle in Figure 5a. The second principal component of these data (PC2) is by definition orthogonal to

Figure 5.

Analysis of simulated experiment. **a:** A scatter plot of the hypothesized fMRI signal values (at times t = 1 and t = 2) for each brain voxel contains arrows IC1 and IC2, which show the directions determined by the relative activations of the two component processes (Fig. 4). The assumption of spatial independence of the IC1 and IC2 maps implies that the data will vary independently along these two component vectors. The two parallelograms (solid borders) indicate the active voxels for each component (e.g., the voxels highlighted in Fig. 4b). Active voxels by correlation analysis are those whose projections onto the reference function (COR) exceed an arbitrary correlation threshold, e.g., those enclosed by the rectangle (dashed borders). The first principal component of the data set (PC1) points in the direction of maximum variance of the data, but has no direct relationship to the two independent component directions (IC1 and IC2). Active voxels associated with the first principal component are those lying inside the tilted rectangle (dotted borders). ICA, PCA, and correlation analyses find overlapping, but typically not identical, collections of active voxels. Only ICA will find the active areas of each independent component (Fig. 4b). **b:** The independent component directions IC1 and IC2 can be indirectly determined by finding the linear transform W, which results in a rectangular distribution. **c:** The sigmoid transformation g(WX) produces the most uniform (i.e., maximum entropy) distribution for the data shown. The ICA algorithm of Bell and Sejnowski [1995] adjusts IC1′ and IC2′ to maximize the entropy of the distribution.

the first component, but also has no particular relationship to either of the independent components.

Figure 5a shows that there may be overlap in the collections of active voxels determined by correlation analysis, PCA, and ICA, but these three methods for finding voxels activated during an experimental task usually will not give identical results. To the extent that assumptions of linear summation, spatial sparsity, and statistical independence between components are valid, ICA should more accurately determine the exact spatial extents and time courses of task-related as well as nontask-related activations contributing to the data.

METHODS

Subjects and image acquisition

A total of 4 normal volunteer subjects participated in two fMRI experiments. In the first, 3 subjects performed a Stroop color-naming task. In the second, a fourth subject performed a word/number task. Each experiment consisted of two 6-min trials of the same task, interspersed by trials involving other cognitive tasks, not reported here. Each trial consisted of five 40-sec control blocks alternating with four 40-sec experimental task blocks.

Subjects' BOLD signal brain activity was scanned using a 1.5T GE Sigma MRI system GE Medical Systems (Waukesha, WI) equipped with an inserted three-axis balanced torque head gradient coil designed for rapid switching [Wong et al., 1992]. A midsagittal localizer slice assisted in determining landmarks for 8–10 (5 mm thick, 1 mm interslice gap) 64 × 64 echoplanar, gradient-recalled (TR = 2500 msec, TE = 40 msec) axial images with a 24-cm field of view. For each slice, 135–146 images were collected at 2.5-sec sampling intervals. Slices were selected to include the anterior cingulate gyrus, implicated by PET studies in Stroop performance [Bench et al., 1993], and portions of the parietal, occipital, and temporal lobes. High-definition anatomical images were also acquired, using a spoiled GRASS protocol to define the localization of the BOLD signal changes with respect to brain anatomy.

Tasks

Stroop task

The Stroop color-naming task is often used to examine disinhibition and selective attention deficits in

Figure 6.
Time smoothing. Illustration of the technique used for temporal smoothing and time alignment. A three-point Hanning smoothing filter was convolved with the data, using slightly different lags for each brain slice to minimize offsets introduced by the successive 250-msec sampling delays in the multislice acquisition process.

patients with brain disorders [Lezak, 1995]. Stimuli spanning a visual angle of 2° by 3° were presented one at a time by overhead projector onto a screen placed at the foot of the magnet. Subjects viewed this screen through a mirror attached to the head coil. A personal computer containing a Cognitive Testing System (Digitry, Inc., Edgecomb, MA) controlled stimulus presentation. Stimuli were presented as near as possible to the center of the subject's visual field. In all conditions, subjects were instructed to covertly name the color of each stimulus, which was red, green, or blue. In control blocks, the subjects were simply required to covertly name the color of a displayed rectangle. During experimental Stroop-task blocks, subjects were required to name the color of the script used to print one of the same color names (i.e., "red," "green," or "blue"). Each color name was displayed in a different color from the one it was named. For example, when the word "red" was presented in blue script, the subject was to think (but not speak) the word "blue." Each trial comprised four task cycles, each consisting of a 40-sec control block and a 40-sec experimental block, followed by a final 40-sec control block. The first 6-min trial was repeated about 15 min after its initial presentation (i.e., after two similar intervening trials). Each subject was pretested during a training session to determine the inter-item interval for which they would make verbal errors on 10–20% of the presented items. This interval was then used in the experiment.

Word/number task

The Brown and Peterson word/number task [Peterson and Peterson, 1959] has been used in experimental

psychology and neuropsychology to investigate word-forgetting over brief intervals [Lezak, 1995]. Stimuli spanning a visual angle of 1° by 2° were presented via the same apparatus as described for the Stroop task trials. During control blocks, the subject simply fixated on an asterisk displayed in the screen center. In the word/number task blocks, the subject passively observed a word that was displayed for 2 sec. During the following 6 sec, an integer between 100–900 was shown on the screen, and the subject was to mentally add successive 7s to it while still remembering the word. For example, if the number displayed were 300, the subject was to think covertly, "307, 314, 321. . ." The subject was not asked to explicitly recall the presented word. Each 40-sec task block contained five word/number stimulus pairs.

Preprocessing

A set of 8–10 slices was collected in cyclic order every 2.5 sec. This rate is faster than the time constant


Figure 7.
Relative contributions of ICA and PCA components. The eight upper traces show the fractional contributions to the observed data of the 144 ICA components for each of the eight trials, rank-ordered by contribution size. The rank of the consistently task-related ICA component in each trial is indicated. These distributions of relative component contributions were highly similar across trials, and differed from the distribution of rank-ordered contributions of the PCA components to the data from one of the trials (bottom trace). As expected, PCA accounted for much of the data variance by a few large components, whereas the relative contributions of the ICA components, specifying the spatially independent components comprising the signal, were more equal.
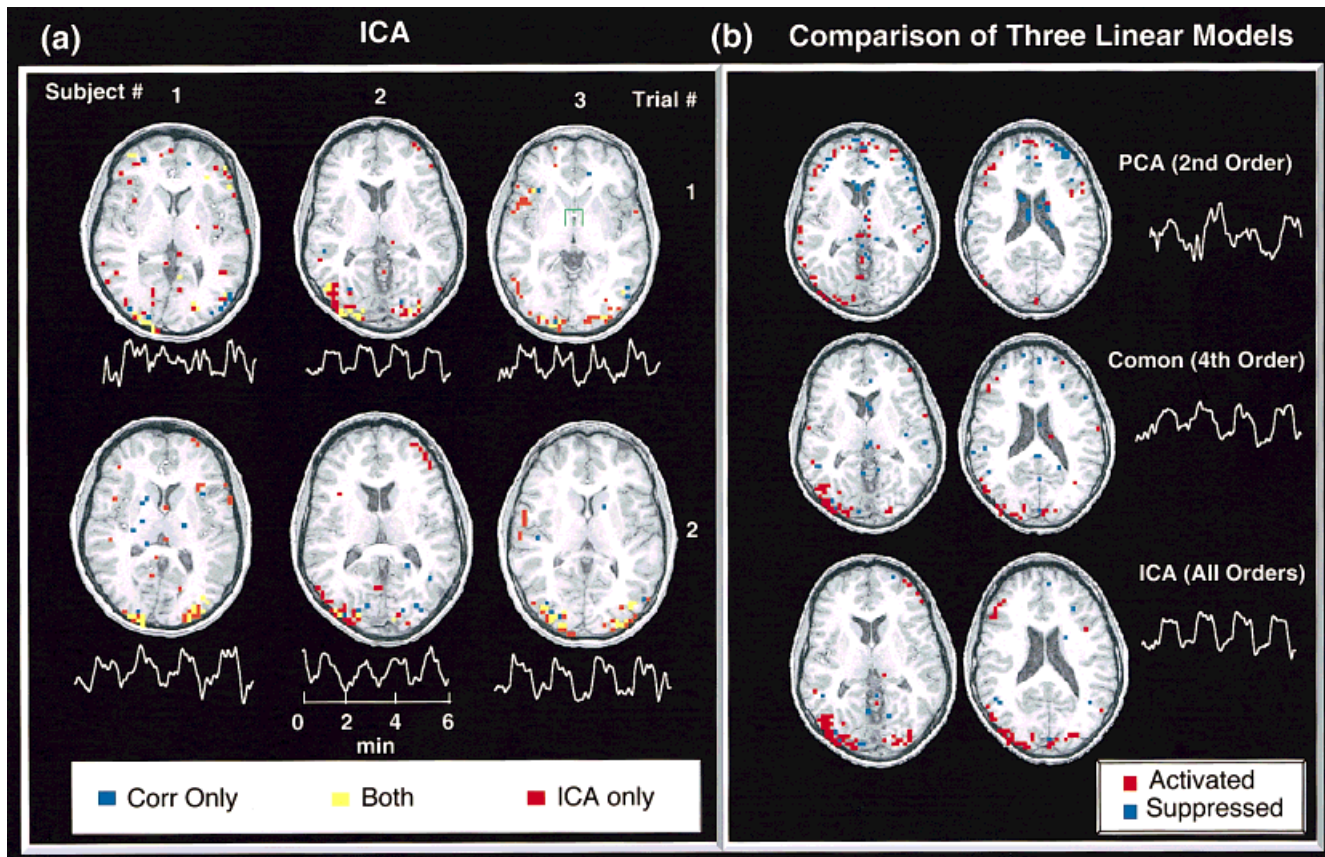
Figure 8.
Consistently task-related (CTR) components. (a) Results for the 3 subjects performing the Stroop color-naming task. Each subject participated in two 6-min trials composed of alternating 40-sec control and Stroop task blocks. ICA decomposition of each trial produced one component whose time course of activation strongly resembled the task block structure (r = 0.64–0.94). Active voxels ICA component map voxels ($|z| > 2.0$) are shown in red, together with voxels considered active by correlation in blue (r > 0.4) for one brain slice. Voxels deemed active by both methods are shown in yellow. Dorsolateral frontotemporal activations were detected only by ICA. (b) Comparison of CTR component maps for PCA, the fourth-order technique of Comon [1994], and ICA. The most consistently task-related component maps are shown for one of the Stroop sessions (subject 2, trial 1) (cf. Fig. 9). Axial slices reveal more focal regions of activity by ICA and the fourth-order decomposition of Comon [1994], while the PCA map is more speckled or diffuse, and does not reveal the extensive occipital activations shown by the other decompositions as well as by correlation (a). Red voxels are activated with the shown time course, while blue voxels are activated *opposite* to the shown time course.

of the BOLD signal hemodynamic response function that typically peaks 5–8 sec after stimulus onset [Bandettini et al., 1992]. The data were not registered to correct for head movement. Voxels indexing active brain regions were determined by examining the mean value of the time series of each voxel. These mean voxel values invariably followed a bimodal probability distribution. The local minimum between the two peaks of a third-order polynomial fitted to the voxel mean-value histogram determined a cutoff value above which voxels were assumed to contain active brain signal. Voxels with weaker signals were found to lie almost exclusively outside the brain and were therefore excluded from analysis.

As shown in Figure 6, data were temporally smoothed using a 3-point filter based on a Hanning window [Press et al., 1992]. The three points were shifted along the window by 250 msec for each successive slice, to decrease the time misalignments induced by the 250-msec acquisition delays between slices. The filtered BOLD signals from all brain voxels at each time point were placed into subsequent rows of the data matrix. The mean of each row (time point) was then subtracted from the data.

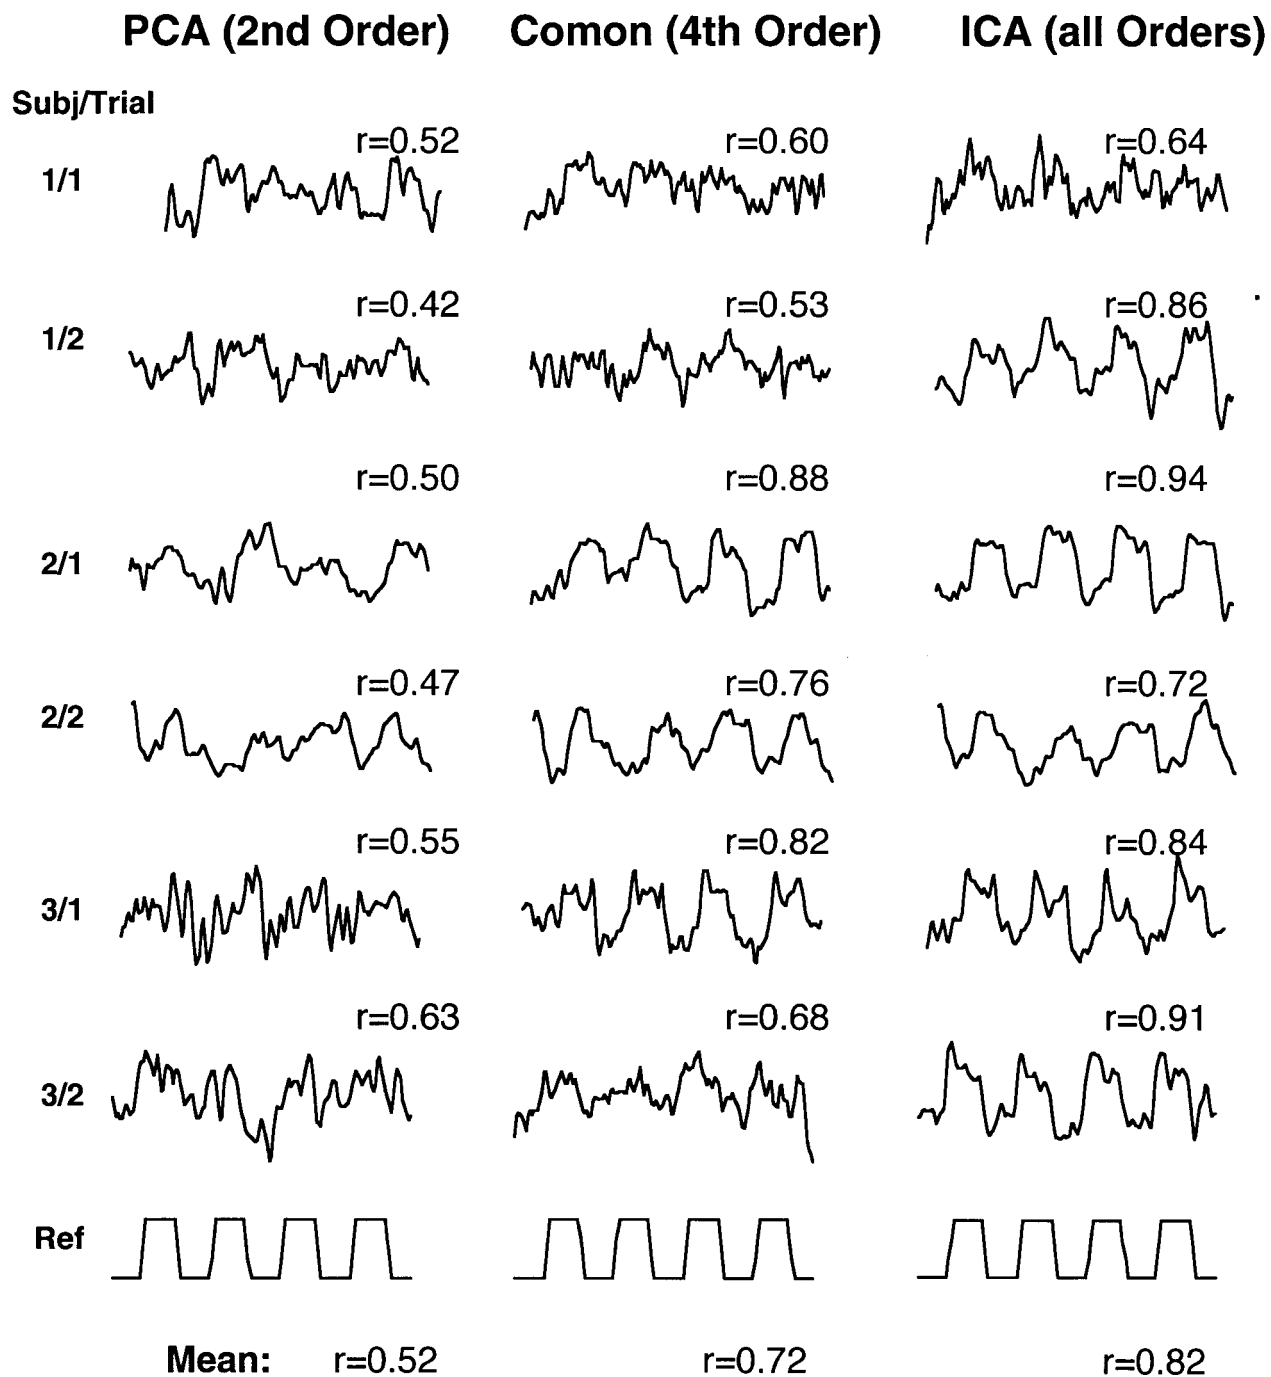# Consistently Task-related
Components from Three Linear Models

PCA (2nd Order)	Comon (4th Order)	ICA (all Orders)

Subj/Trial

1/1	r=0.52	r=0.60	r=0.64
1/2	r=0.42	r=0.53	r=0.86
2/1	r=0.50	r=0.88	r=0.94
2/2	r=0.47	r=0.76	r=0.72
3/1	r=0.55	r=0.82	r=0.84
3/2	r=0.63	r=0.68	r=0.91

Ref

Mean: r=0.52   r=0.72   r=0.82

Figure 9.

Component independence. Comparisons of three linear decomposition techniques, PCA, the fourth-order algorithm of Comon [1994], and the ICA algorithm of Bell and Sejnowski [1995]. For each of the three techniques, the component time course most closely matching the Stroop task reference function is shown. As the imposed criterion for spatial independence between maps becomes more strict, from PCA (second order), to the technique of Comon (fourth order), to the current ICA method (all orders), there is stronger agreement between the CTR-component time course and the reference function.

The ICA algorithm was applied separately to data from two 6-min trials for each subject. Analysis was performed using a matrix code implemented in MAT-LAB 4.2 (Mathworks, Inc.). Convergence of the ICA analysis for each 6-min trial session typically took 90 min on a Digital Equipment Corporation Alpha 2100A computational server.

Once W had been determined by the algorithm, component maps C were derived using Equation (4). The time course of activation of each component was contained in the corresponding column of W^{-1}. For comparison, the eigenimages from each trial were determined using standard PCA techniques [Jackson, 1991], along with their associated time courses. To determine the effects of higher-order statistics on determining uncorrelated spatial maps, an ICA technique using fourth-order cumulants proposed by Comon [1994] was also used to find partially independent maps and associated time courses. Convergence of the Comon algorithm typically took 360 min of computer time per trial.

For each trial, a reference function was constructed by convolving a square wave matching the time course of the experimental/control task blocks with a crude approximation of the BOLD impulse response function, a rectangular function of 7.5-sec duration (Fig. 1). This reference function was then correlated with the signal time course of each voxel [Bandettini et al., 1993] and with the time courses of the maps derived by the ICA and PCA techniques

$$r_k = \frac{\sum_{i=1}^{n} (x_{ik} - \bar{x}_k)(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n} (x_{ik} - \bar{x}_k)^2} \sqrt{\sum_{i=1}^{n} (y_i - \bar{y})^2}} \qquad (17)$$

where r_k is the correlation coefficient for the k^{th} voxel, x_{ik} is the recorded value of the k^{th} voxel at the i^{th} time point, and y_i is the reference function at the i^{th} time point.

A cutoff value of $r_k = 0.4$ was used as the correlation significance threshold. Voxels whose r_k exceeded this limit were considered *active voxels* by the correlation method. ICA revealed that one trial from one subject contained a prominent linear trend. Therefore, the data from this trial were linearly detrended before correlation analysis for fair comparison with the ICA results.

Here,

$$X'_{ij} = X_{ij} - [m_j i + b_j] \qquad (18)$$

where X_{ij} is the recorded value of the j^{th} voxel at the i^{th} time point, x'_j is the time series of the j^{th} voxel after detrending, and m_j and b_j are defined by

$$m_j = \frac{n\left(\sum_{k=1}^{n} k\, X_{kj}\right)\left(\sum_{k=1}^{n} k\right)\left(\sum_{k=1}^{n} X_{kj}\right)}{n\left(\sum_{k=1}^{n} k^2\right) - \left(\sum_{k=1}^{n} k\right)^2} \qquad (19)$$

$$b_j = \frac{\left(\sum_{k=1}^{n} X_{kj}\right)\left(\sum_{k=1}^{n} k^2\right) - \left(\sum_{k=1}^{n} k\right)\left(\sum_{k=1}^{n} k\, X_{kj}\right)}{n\left(\sum_{k=1}^{n} k^2\right) - \left(\sum_{k=1}^{n} k\right)^2} \qquad (20)$$

and n is the number of time points in the trial.

For each trial, the computed ICA component and correlation active voxels were read into the functional neuroimaging display program MCW AFNI [Cox, 1996] for display and registration with the structural T1-weighted MRI brain images.

RESULTS

ICA, PCA, the fourth-order method, and correlation analysis were applied to the eight data sets for the Stroop and name/word task outlined in Methods. The

Figure 10.
Consistent activation across trials during the Stroop task. The CTR component maps from both trials from each subject were spatially smoothed with a 3-D, 6-mm, full-width half-maximum Gaussian kernel and averaged over trials. The scatter plots (at right) plot the smoothed voxel z-values from the CTR map in trial 2 (axis: $-5 < z < 5$) vs. the smoothed z-values in the CTR map obtained from trial 1 (axis: $-5 < z < 5$), along with the correlation of each voxel with the reference function in trial 2 (axis: $-0.5 < r < 0.5$) vs. the correlation values in trial 1 (axis: $-0.5 < r < 0.5$). The oblique lines in the scatter plots correspond to the z-value thresholds labeled in the ICA CTR component maps. Correlational thresholds were those giving an equal number of active voxels as the number of active voxels in the ICA CTR component at the various z-thresholds. Note that frontal activity in the second subject was detected only by ICA (middle).

Figure 10.

contributions, γ_i, of each ICA component to the data were computed as in Equation (6). These ranged from $0.08–5 \times 10^{-4}$. The distributions of the ICA component contributions were similar across trials and were quite unlike the distribution of contributions of projections on the principal components of the same data (Fig. 7). Some maps contained multifocal groupings of active voxels, while others (usually those with contributions < 0.01) had diffuse or "speckled" spatial distributions. The time courses of the components could be grouped into broad classes. The time courses of some components followed part of or the entire task block design (CECECE...), while others were slowly varying, quasiperiodic, or noisy in appearance.

Consistently task-related components

In all trials, exactly one ICA component had a time course that was highly correlated (r = 0.64–0.94) with the reference function. In all cases, this *consistently task-related* (CTR) component had a relatively low contribution rank (14th–41st). Maps of active voxels for these task-related components (using a $|z| > 2$ threshold) contained areas of activation resembling those produced by the correlation method with r > 0.4. During Stroop trials, ICA and the correlation method detected task-related activation in Brodmann's areas 18 and 19 (not involving the calcarine fissure) and in the supplementary motor area and cingulate system. In each of the trials, the ICA method also detected task-related activation in prefrontal areas, including the left dorsolateral prefrontal cortex (Figs. 8, 10, 14, 17). The first subject performing the Stroop task was later found to have been vocalizing the words during the first trial rather than stating them covertly, probably introducing additional motion artifact in the data.

Figure 9 compares the time courses of the components best matching the reference function for each of the three linear models used in the Stroop task: PCA, the fourth-order method, and the ICA algorithm described. Several PCA component maps (Fig. 8b) had associated time courses that were moderately correlated with the reference function, although these correlations were lower than for the CTR ICA components. In general, as successively stricter criteria for spatial independence between individual maps were applied, the time courses of the maps more closely matched the reference function (Fig. 9).

In order to detect areas of activation consistent across trials, the CTR and correlation maps from each of the two Stroop trials from each subject were aver-

aged after spatially smoothing each map with a three-dimensional (3-D), 6-mm, full-width-half-maximum Gaussian kernel. As shown in Figure 10, the frontal activation detected in all 3 subjects by ICA was robust to changes in the z-threshold for activation. Reducing the threshold added active areas adjacent to regions of activation found with higher thresholds. Only the ICA method detected frontal activation in the second subject, even after a significant reduction of the correlational threshold (r = 0.16). Frontal activation was detectable by correlation in the third subject, but only by reducing the threshold (r = 0.23) until apparent activation in the basal ganglia, thalamus, and lateral ventricles was observed.

In Figure 11, the four 80-sec task cycles of the CTR ICA component in each of the Stroop trials are superimposed. In each trial, the fine temporal structure of the activation was stereotyped within subjects. The right side of Figure 11 shows the mean of the eight ICA component task activations in the two trials from each subject, superimposed on the expected response (one cycle of the task reference function). Note that the mean time courses for each subject (Fig. 11, right column) were not reliably estimated by the reference function, suggesting that the true extent of hemodynamic activation during Stroop task performance was not constant but tended to decline during the course of the experimental blocks. All 3 subjects showed greater activation near the beginning of the trial. Subjects also differed in the rise-time of activation. These details tended to be consistent across task cycles. Further, the time course given by ICA much more closely resembled the mean time courses of the most active voxels, as determined by either ICA or correlation, than did the idealized reference function.

For the subject performing the word/number task, areas of significant activation by ICA and correlation were again similar (Fig. 12). Both methods found task-related activation in Brodmann's areas 18 and 19 and in left occipital-parietal areas. ICA also indicated significant activation in frontal and temporal regions.

Figure 13 shows a scatter plot displaying each voxel's value in the consistently task-related map vs. its correlation with the reference function for one Stroop trial (subject 2, trial 1). Both methods found 47 active voxels in common (Fig. 13, upper right). ICA also found 175 voxels whose correlation with the reference function was <0.4 (including some whose correlation with the reference function was near zero). However, the mean time course of these 175 voxels (Fig. 13a, upper middle) clearly reflected the alternating task-block sequence, supporting the implication of
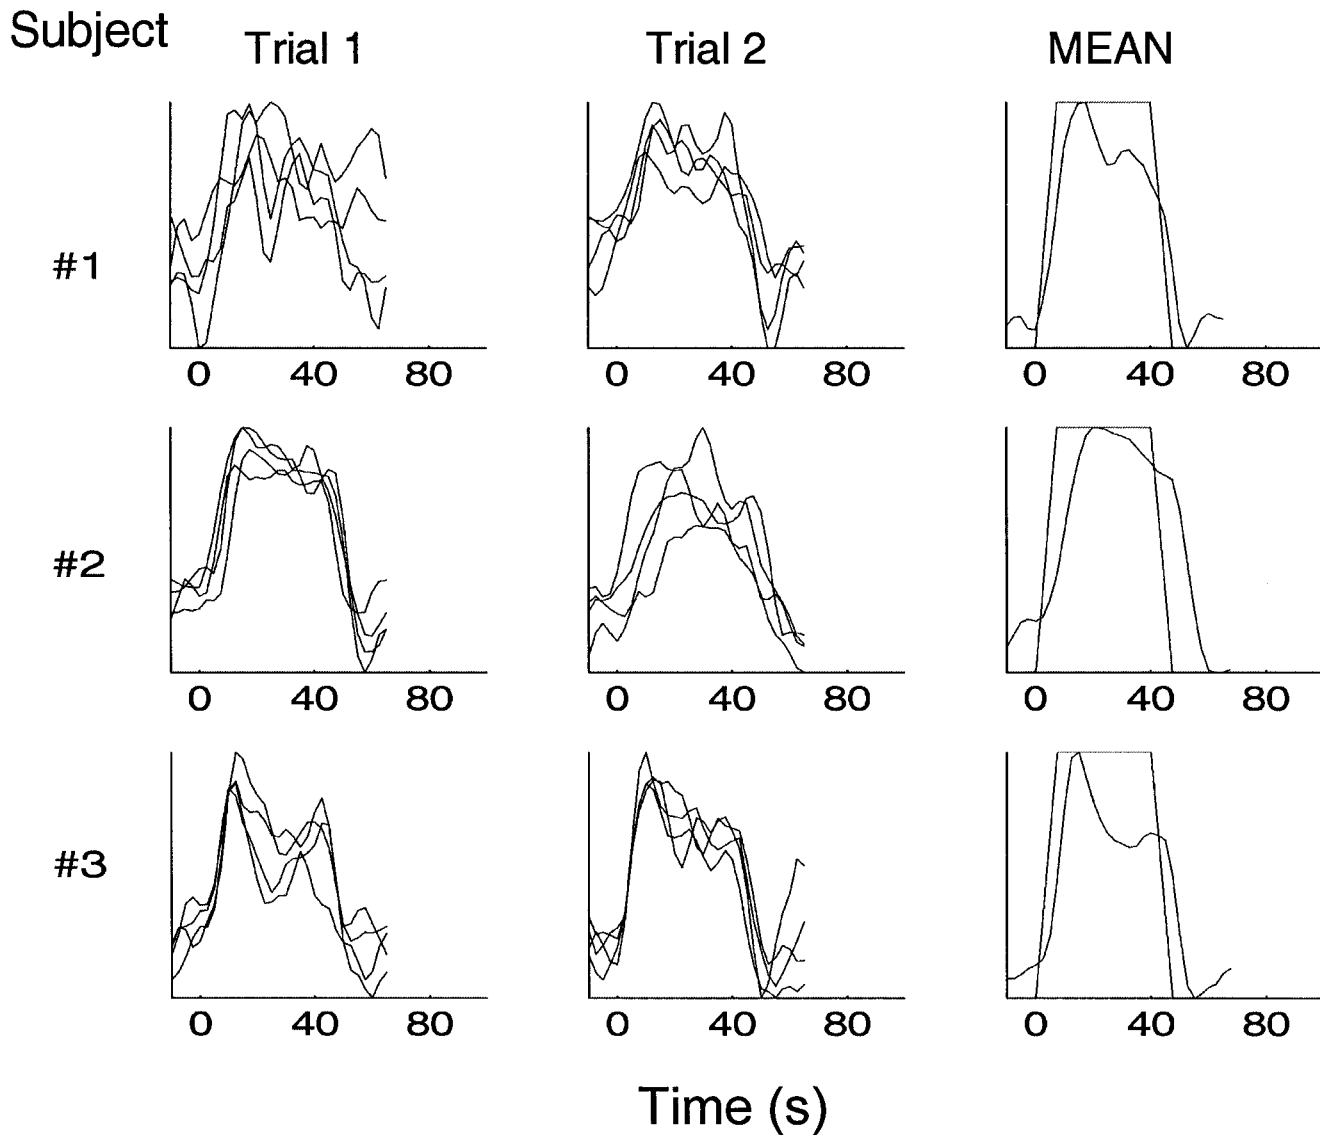
Figure 11.

Consistency of ICA in task-related activations. At left and center are superimposed the four successive 80-sec task cycles of the consistently task-related component activations from each of the six Stroop trials. Right: The means of the eight task-cycle activations for each of the 3 subjects. A single cycle of the reference function used in the correlation analysis is superimposed on the component means for comparison. Note the stereotyped details of the experimental task activations in at least four of the trials, and the individual subject differences between the mean activations and the assumed task reference function.

the ICA results that activity at these voxels was influenced by task performance.

Since the data for this trial contained a prominent linear trend, linearly detrending the data before correlating with the reference function (Fig. 13b) increased the overlap between the results of the two techniques (from 47 voxels in common to 105). However, there were roughly as many voxels that each method detected individually (117, 113) that the other did not.

After linear detrending prior to correlation, frontal activation was detected by both methods (Fig. 14). The failure of ICA to detect as significant some of the active voxels detected by correlation might be explained by the participation of these voxels in other ICA components *transiently* time-locked to the task block design (see below) or by inaccuracy in selecting equivalent ICA and correlation thresholds (here z = 2.0, r = 0.4). The mean time course of the 105 active voxels by both

Figure 12.
Word/number task activations. Consistently task-related component activations for the subject performing the word/number task in two trials. Again, the ICA decomposition included a single component whose time course was highly correlated with the task block structure of the trial. This component had more active voxels (z > 2.0) in posterior visual association areas than were found by correlation with the reference function (r > 0.4). ICA also found active frontal and lateral regions not detected by correlation analysis.
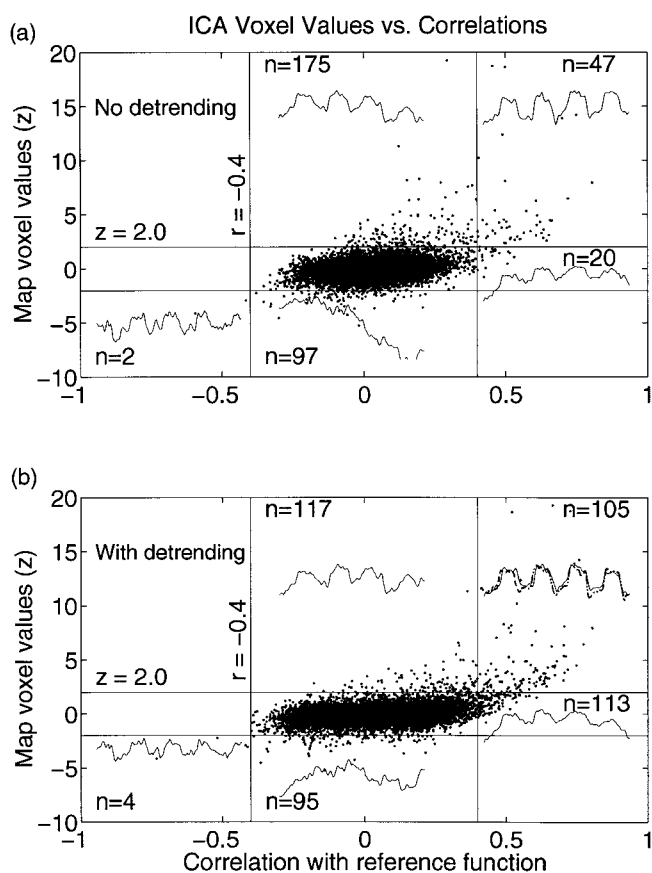
Figure 13.

ICA vs. correlation. **a:** A scatter plot comparing voxel values in the CTR component map (subject 2, trial 1) to correlations of the signal at each voxel with the task reference function. Horizontal lines ($z = \pm 2$) separate the voxels into active and inactive subsets according to ICA, while vertical lines ($r = \pm 0.4$) indicate the threshold for active voxels used in the correlation analysis. The numbers of voxels falling in the resulting portions of the plot are noted. The plotted waveforms represent the mean time courses of the voxels in each portion. Forty-seven voxels were considered active by both analytical methods (upper right). The ICA method selected 175 voxels as active that were considered inactive by correlation (upper center). The mean time course of these voxels clearly showed task-related activation (upper center trace). Correlation marked 20 voxels as active whose ICA map values were considered inactive (right center). **b:** When the linear trend was removed from the time course of each voxel before correlating with the reference function, the number of voxels considered active by both methods increased to 105 (top right), and the numbers of voxels considered active by one method only (117/ 113) were equalized. Note that the mean time course of the 105 voxels detected by both methods (solid line, upper right) was highly similar to the detrended CTR component time course (broken line, upper right).

methods (solid trace, upper right) closely resembled the CTR component time course (broken trace).

The unique consistently task-related component in each trial had a multifocal character, as shown in Figures 8a, 12. The other 140 or more components for each trial could be grouped empirically into several broad classes, described below according to general features of their spatiotemporal structure.

Transiently task-related components

Some components appeared to be time-locked to the task-block design during part of the trial. For example, the active areas for the component shown in Figure 15a included frontal and occipital regions. This component was abruptly activated during the second Stroop task block but not during other task blocks. Such transiently task-related (TTR) activity might not be detected by a correlational analysis that averaged over all the task-block cycles in a trial.

Slowly varying components

In most trials, there were also slowly varying components (Fig. 15b). In the trial shown, voxels indexing regions of the ventricular system were separated into one slowly-varying component (Fig. 15b, solid line), implying that part of the BOLD signals at these voxels changed in synchrony with the time course of this component. The dotted line in Figure 15b shows the mean time course of the active voxels for this component. Note that the time courses of the components shown, although monotonic, are not linear and therefore could not be removed entirely by linearly detrending the data.

Quasiperiodic components

In data from each of the 4 subjects, several components had approximately oscillating time courses with bimodal periods near 14 or 40 sec (Fig. 15c). These components showed similar areas of activation in both trials, mostly restricted to a single brain slice.

Movement-related components

Some components had abrupt changes in their time course and/or ring-like spatial distributions, suggesting sudden or gradual head movements. The distribution of positive and negative voxel values for the component shown in Figure 15d and its abruptly shifting time course suggest the effect of a torsional head movement in the coronal plane. Other components had a "ring-like" spatial structure, like those shown

Figure 14.
ICA vs. correlation with linear detrending. The spatial distribution of voxels detected by ICA (red, $z \geq 2$) and by correlation (blue, $r \geq 0.4$) after detrending. Same data set as Figure 13. Note the frontal regions of activation detected by both methods (circles).

in Figure 15e, which we suspect represented motion in the axial plane. Head-movement simulations (reported below) tended to support this hypothesis.

Residual noise components

The smallest ICA components (especially those with contribution rankings of 90–144) had diffuse or "noisy" spatial and temporal patterns and most probably represented noise in the data. Their time courses and maps were not reproducible across applications of the algorithm, even on the same data. The noisy character of one of these small components is clear in Figure 15f, which shows the random distribution of active voxels in two axial slices.

Voxels active in several components

Figure 16 demonstrates that a single voxel could participate significantly in several ICA components of more than one of the types listed above. In Figure 16, the time course of the BOLD signal of the voxel highlighted in the center image is shown beneath. This voxel was highly weighted ($z = 5.0$) in the CTR component (Fig. 16, middle right), but was also active in three other components of various types. The ICA method is able to determine that a voxel is an active participant in a CTR component even though, because of the influence of other component processes, that voxel's time course may not appear to be task-related. Calculations showed that most voxels were active ($|z| > 2$) in 1–6 components, and that on average each voxel was active in 3.16 components.

The spatiotemporal structure of task-related activations

Figure 17 shows four consistently or transiently task-related components from one Stroop trial (subject

3, trial 2). The time courses of the TTR components appear time-locked to that of the CTR component for part of the trial. By summing the contributions of these components, a more complete and dynamic representation of the spatial and temporal structure of task-related activity can be reconstructed.

Tests of reliability and simulations

To further our understanding of the reliability of task-related ICA components, the data set from one trial (subject 2, trial 1) was manipulated in different ways and the resultant data sets were analyzed again by ICA to determine the effects of the manipulations on the computed ICA CTR component.

Robustness against added noise

The standard deviations of the BOLD signal over the course of the selected trial were computed for each of the voxels and then ordered by relative size. The voxels with the smallest signals (those in the lowest quartile) had quite similar standard deviations. Their mean was used as an estimate of the baseline noise level in the data. Independent zero-mean Gaussian noise samples drawn from distributions whose standard deviations were various percentages of the baseline noise were added to *each* time point of *every* voxel in the trial. These noise-added data sets were then analyzed as if they were raw data to determine whether a CTR ICA component could still be extracted. Even after adding Gaussian noise at a level equal to that of the estimated noise baseline, a CTR component was recovered (Fig. 18). Although the exact morphology of the square-wave time course varied slightly between the various runs, correlations between the original CTR component time course and the time courses of the CTR component extracted from the noise-added data were all above r > 0.8.

Component reliability

To further test the reliability of CTR components, the data from one Stroop trial were split into odd and even time points, and each of the two 72-time-point data sets were decomposed separately using ICA. Figure 19 shows that both the odd and even decompositions returned a component whose time course and map voxel value distribution were highly related to those of the original CTR component.

Detection of simulated head movement

The ability of the ICA method to detect abrupt head movement was investigated by artificially simulating a head shift by one voxel in a diagonal direction (4.2 mm) midway through one trial. The largest component of the ICA decomposition of the resulting data consisted of a step function at the appropriate time point (Fig. 20a). The map for this component was concentrated at the cortical surface with opposite signs in the right frontal and left occipital areas, i.e., the regions of maximum signal change following the simulated movement. Another simulation (not shown) demonstrated that the ICA algorithm could also be used to readily detect simulated head movements of one quarter of a voxel (~1 mm). The ICA decomposition of some trials included components whose waveforms contained sharp transient shifts (Fig. 15d) and/or whose maps had a similar ring-like structure (Fig. 15e). We tentatively interpret these components as arising from small abrupt or gradual head movements during the trial.

Detection of a simulated task-related activation

A simulated task-related signal with a three-cycle time course and a spatial distribution unlike that of the actual CTR component was added to (or subtracted from) the signals of voxels in four arbitrarily selected regions of two brain slices (Fig. 20b, top). The level of simulated activation was equal to that of the CTR component of the same trial (Fig. 20b, upper left). ICA was used to decompose the resulting data set. The time course of one of the resulting independent components corresponded to the simulated three-cycle activation, while a second component accounted for the actual four-cycle task-related activation (Fig. 20b, middle panel). Active areas ($|z| > 2$) of the simulated three-cycle component included every voxel in the simulated active regions, plus just two extraneous voxels. Correlation of each voxel time course with the simulated three-cycle reference function, on the other hand, marked as active (r > 0.4) only a small proportion of the affected voxels (Fig. 20b, middle panel).

To test whether the relative insensitivity of correlational analysis in this instance depended upon the selection of too high a correlation threshold, we reduced the threshold until the number of active voxels detected by correlation equaled the number detected by ICA (at $|r| > 0.197$). The reduced correlation threshold produced the active voxel map shown in Figure 20b, bottom panel. In this map, 66% of the active voxels were not in the area of simulated activation. Thus, ICA proved both more sensitive and more specific than correlation in detecting the areas of simulated task-related signal. In addition, the correlation method required that the experimenter knew the time course of

(a) Transiently task-related

(b) Slowly-varying

(c) Quasi-periodic

(d) Abrupt head movement

(e) Slow head movement

■ Activated
■ Suppressed

(f) Noise

Figure 15.

the simulated task activation, whereas ICA determined and separated the approximate time courses of both task-related activations without any a priori information about their possible time courses.

DISCUSSION

Our results indicate that independent component analysis (ICA) can be used to reliably separate fMRI data sets into meaningful constituent components, including consistently and transiently task-related physiological changes, nontask-related physiological phenomena, and machine or movement artifacts. For the ICA method to separate task-related activity from other component activity, the spatial distribution of brain areas activated by task performance must be spatially independent of the distributions of areas affected by artifact. Confidence in this assumption is strengthened by the result that the time courses of the CTR components in six Stroop task trials more clearly resembled the block design, as successively stricter criteria for spatial independence were applied to linear decompositions of the data (Figs. 8, 9). The time course of the CTR component determined by PCA did not resemble the task block design as well as the CTR components from the fourth-order or higher-order ICA algorithm described here. The algorithm of Bell and Sejnowski [1995] is computationally more efficient than the technique of Comon [1994] and converges to quite similar decompositions, independent of the initial weights and random seed used in the training [Makeig et al., 1997]. Our simulations (Fig. 18) indicate that the results are robust in the presence of noise in the data. Furthermore, as the ICA model gives a linear decomposition of the data, its results are easy to manipulate. Even for experiments with a simple repeated block design, the ICA method appears to be more sensitive in detecting task-related activation than correlating with an idealized reference function. Thus, we detected variable frontal activation in the CTR component of all 3 subjects performing the Stroop task, which was in some cases undetectable by correlation methods (Fig. 8). Several transiently task-related components also demonstrated dorsolateral prefrontal activity (Fig. 17), which would be unlikely to be detected by correlational methods because their time courses could not be known in advance. Variable frontal activation during sustained or repeated task performance has been reported in several previous PET and fMRI studies, and may relate to changes in subject visual-spatial attention [Nobre et al., 1997], language processing [Binder, 1997], changes in stimulus novelty [Tulving et al., 1996], verbal fluency [Phelps et al., 1997], verbal suppression [Nathaniel-James et al., 1997], and working memory [Manoach et al., 1997], all of which may be required for repeated Stroop task performance in either normal or impaired subjects.

ICA also produced quasiperiodic components many with periods between 10–20 sec (Fig. 15c). As it is not possible to bandpass-limit BOLD signals prior to data collection, periodic signal changes faster than 0.2 sec/cycle (for the given sampling interval of 2.5 sec) are "aliased" back into the captured signal, and may appear in any frequency range. Quasiperiodic fMRI signal fluctuations, therefore, might be caused by aliased cardiac (~1/sec) and respiratory ($\sim1/4$ sec) rhythms [Biswal et al., 1996; Kwong, 1995; Le and Hu, 1996; Mitra et al., 1997]. Much slower cerebrovascular waves, presumed to be due to autoregulatory feedback of cerebrovasculature [Chichibu et al., 1995; Wayenberg et al., 1995], may also be a potential mechanism for producing fluctuating BOLD signal changes. Transmitted pulsatile movements may also precipitate a BOLD signal response throughout the whole brain via induced pressure changes, while blood flow effects are usually local to the great vessels [Kwong, 1995].

The single-slice appearance of the quasiperiodic components in all trials might possibly also reflect the

Figure 15.
Other ICA component types. Types of components detected by ICA decomposition (red, $z \geq 2.0$; blue, $z \leq -2.0$). Negative z-values mean that those voxels are activated *opposite* to the plotted time course. (a) Transiently task-related (TTR) component (subject 1, trial 1, rank 33). This component was selectively activated during the second experimental block. The dotted line shows the time course of the consistently task-related component for comparison. (b) Slowly varying, nontask-related component (subject 1, trial 1, rank 12). The active region for this component was the ventricular system. The lower trace (dotted line) shows the mean time course of the active voxels ($z > 2.0$). (c) Quasiperiodic component. This component (subject 1, trial 4, rank 40) was mostly active in a single slice and had a dominant period of about 14 sec. Other similar components were active in other slices. The spatial distributions of such components were highly reproducible between trials. (d) Suspected abrupt head movement (subject 1, trial 4, rank 18). The right-temporal pattern of active voxels for this component could be produced by a small, abrupt torsional head movement. (e) Suspected gradual head movement (subject 1, trial 4, rank 20). The left/right ring-like pattern of active voxels, together with the monotonic time course, suggests a slow head shift. (f) Residual, "noisy" component (subject 1, trial 4, rank 69). Almost all the smallest ICA components were of this type and appear to represent noise in the data.

Component
rank

Map Value

Selected voxel

#6 (z= -5.9)

Observed voxel
time course

#71 (z= -3.2)

#14 (z= -4.0)

#41 (z=5.0)

Figure 16.
Summation of ICA components. Most brain voxels were active (i.e., had map values of $|z| > 2$) in 1–6
ICA components (mean, 3.19). Here, one voxel in a posterior visual association area participates
strongly in the CTR component for this trial (z = 5.0) as well as in two other larger (lower-rank)
components and one smaller (higher-rank) component.

spin-excitation history used in the acquisition of a
brain slice. Attempts have been made to explicitly
model spin-excitation history to counteract this artifact
[Friston et al., 1996]. However, as ICA reliably and
consistently extracted and separated quasiperiodic
components from other components, the ICA tech-
nique might circumvent the need to explicitly model
spin-excitation history for this purpose.

The ICA method assumes that the observed fMRI
data are the linear sum of components with unique
(though possibly correlated) time courses and statisti-
cally independent distributions of map voxel values.
The method can be viewed as a version of the "general
linear model" [Friston, 1996] currently used in func-
tional neuroimaging and given by

$$X = G\beta + e$$

where X is a data matrix with elements x_{ij} (the
observation of the j^{th} voxel at the i^{th} time), G is a
"design matrix" specifying the time courses of all the
factors hypothesized to be present in the observed data
(e.g., the task reference function, or a linear trend), β is
a matrix of map voxel values for each hypothesized
factor, and e is a matrix of noise or residual modeling
errors. Given this linear model and a design matrix G,
the maps β can be found by least squares estimation. In
contrast, the ICA method extracts *intrinsic* spatially-
independent components of the observed data and
determines explicitly their time courses, rather than
relying on a priori hypotheses as to what they should
be. The need for procedures that involve splitting an a
priori design matrix G into parameters of interest and
of no interest, in an attempt to increase signal-to-noise
ratio [Friston, 1996], might thus be reduced by using

Figure 17.
Task-related components for one trial (subject 3, trial 2). The *consistently task-related* (CTR) component map and its associated time course are shown at bottom. Other components had time courses (yellow lines) that appeared time-locked to the CTR for that trial (white lines) during part of the trial (blue rectangles), and so could be called *transiently task-related* (TTR).

the ICA technique. Our results demonstrate that ICA can extract both transiently and consistently task-related, nontask-related, and artifactual components without a priori knowledge of their temporal or spatial structure. This property of the ICA algorithm warrants its description as providing "blind separation" of the data into spatially independent components.

Although the algorithm is capable of "blind separation" into independent components, the subsequent interpretation of the separated components requires additional knowledge on the part of the experimenter. In the current trials, which utilized a task block design, the one component for each trial that appeared to be consistently task-related was easily found by comparing the component time courses to the task reference function. Although we did not know a priori the exact time course of activation, rough knowledge of the task block design was required to identify the appropriate component. We are currently investigating heuristic approaches for objective classification of the separated

components. For example, the ring-like spatial structures of some components are suggestive of head movement, and the erratic temporal and spatial patterns of other components suggest that these may represent noise in the signal.

If one or more of the ICA components derived for a given data set are identified as artifactual, it is possible to reconstruct the data with these components removed by zeroing the appropriate rows of C in Equation (5). This potentially allows ICA to be used as a preprocessing step prior to further analysis by any other technique. However, movement artifacts cannot be totally removed by this method, as the changes in a voxel's signal activity due to encroachment of a neighboring voxel during movement are a violation of the assumption made by ICA that the maps are spatially stationary. Further work is needed to determine how the movement-related components can be used to readjust the raw data to eliminate the detected movement or errors in registration for subsequent reanalysis.
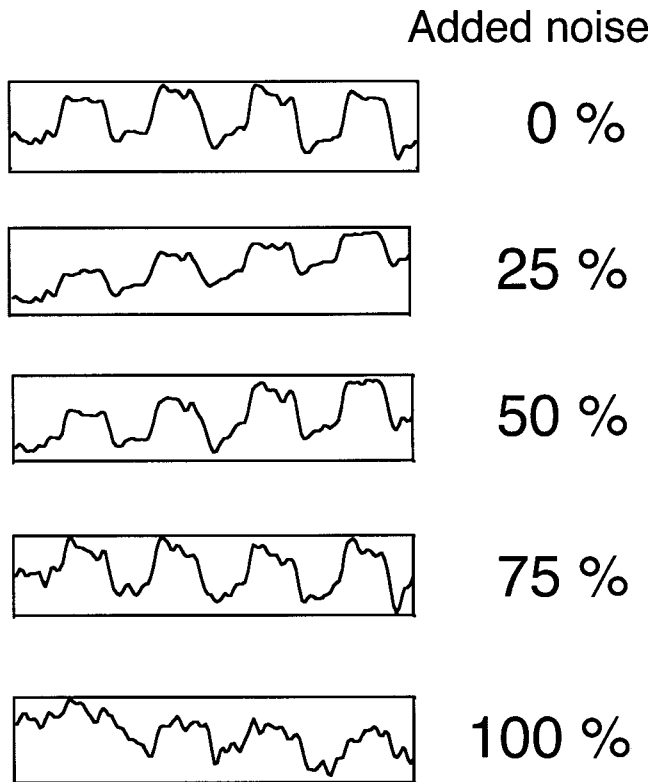
Added noise

0 %

25 %

50 %

75 %

100 %

Figure 18.
Robustness to added noise. The ICA method separated a single CTR component even after the addition of Gaussian noise to each voxel time course at levels 25%, 50%, 75%, and 100% of baseline.

Although the ICA method appears to be useful for fMRI data analysis, it also has some inherent limitations. First, fMRI signal component processes may exhibit saturation or other nonlinear properties and thus may not be appropriate for analysis using a wholly linear model. However, since the task-related components in these and other experiments generally make small contributions to the baseline BOLD signal, an assumption of additivity may be reasonable [Boynton et al., 1996]. Second, the ICA algorithm assumes that the distribution of voxel values specifying the map for each signal component is statistically independent of the distributions of voxel values specifying all the other component maps. Although this criterion provides an essentially *unique* decomposition of the data, it may not necessarily be the *desired* representation for all purposes. The spatial independence criterion, together with the particular (here, logistic) nonlinearity used in the algorithm, biases the ICA method towards finding components having relatively sparse as well as discrete active component areas [McKeown et al., 1998]. If some component process produces proportional signal changes over a relatively large part of the brain, the ICA method used here might split the effects of this process into several ICA components with smaller active areas and closely related time courses. Similarly, if two component processes contribute to the observed fMRI signals in well-overlapping brain areas, ICA may split the resulting activity into three or more components, one component representing the combined effects of the two factors in the regions of overlap, and two others representing the regions affected by just one of the two processes. This may partly explain the numerous TTR components detected by the technique [McKeown et al., 1998]. Nevertheless, the optimum way to describe the varying spatial extent of time-dependent, task-related activations detected in fMRI data is unclear. Combining the task-related (CTR and TTR) ICA components (Fig. 17) may provide a practical, if somewhat cumbersome, method.

One benefit of the ICA technique is the ability to discern activations that could not be predicted in advance of the experiment, e.g., TTR activations. It is possible that TTR components, during times when they are *not* time-locked to the experiment, represent cognitive systems indirectly related to task performance, e.g., arousal or alertness. The ICA approach should also be highly promising for investigations of patients with pathological conditions that may alter the latencies, amplitudes, and brain distributions of their fMRI signals in unpredictable ways.

There are several issues about the ICA decomposition of fMRI data that still need to be addressed: 1) The smallest ICA components, particularly those with speckled spatial distributions, appear to be noise of unknown origin. As yet, we do not know what proportion of a given component is physiological signal or identifiable artifact, and what is noise. 2) The assumption that the component maps are spatially stationary makes the method sensitive to the detection of movement artifact, but does not, in its current form, allow for the straightforward correction of suspected head movements. 3) Methods for testing the statistical reliability of ICA component time courses and areas of activation need to be developed. The ability of the algorithm to converge to equivalent components, using data from a subset of time points from a trial (Fig. 19), suggests that "jackknife" or other bootstrap methods can be employed to determine levels of statistical significance for the voxel values in a map. In such approaches, components computed from training on subsets of time points are compared to estimate the robustness of the statistics derived from the complete data set (e.g., component map values and time courses).
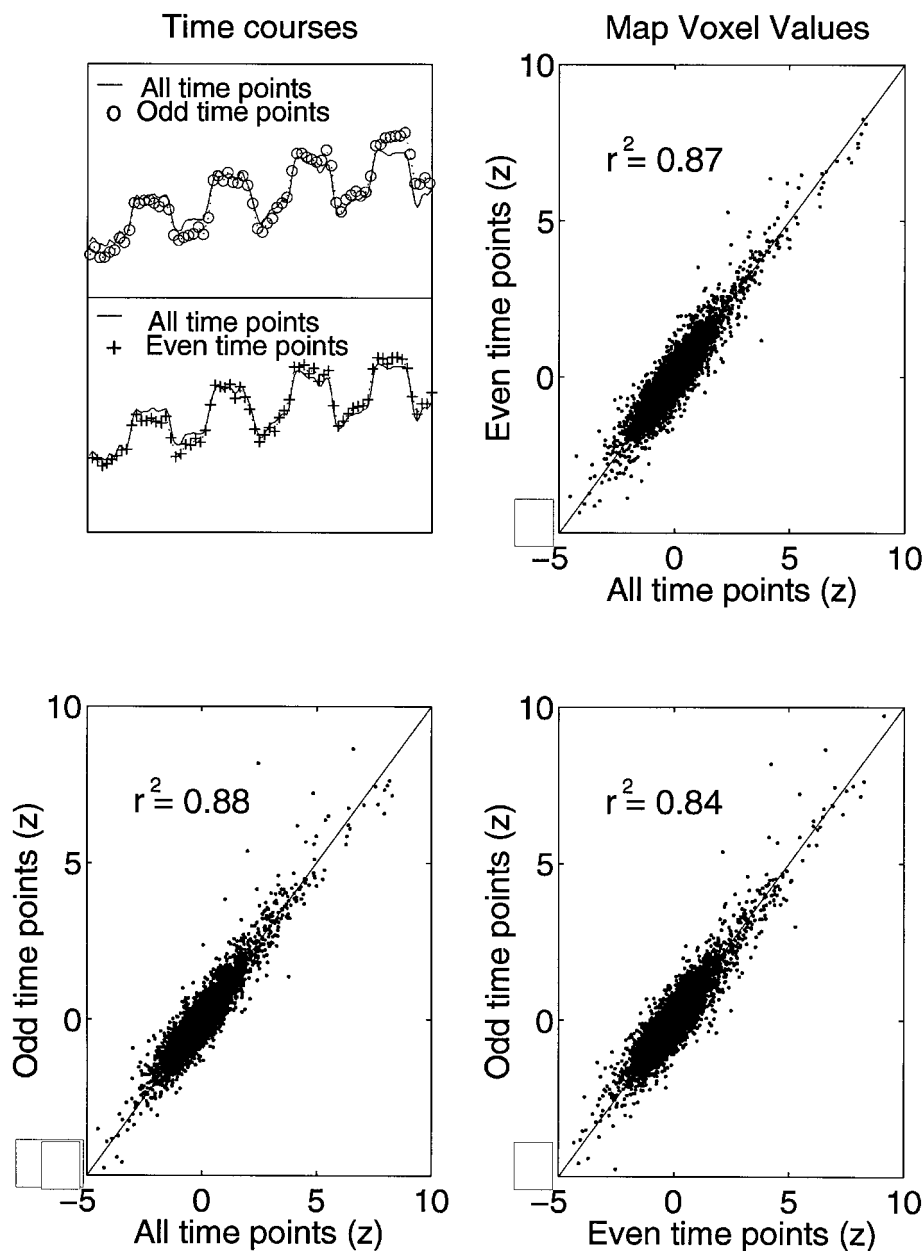
Time courses
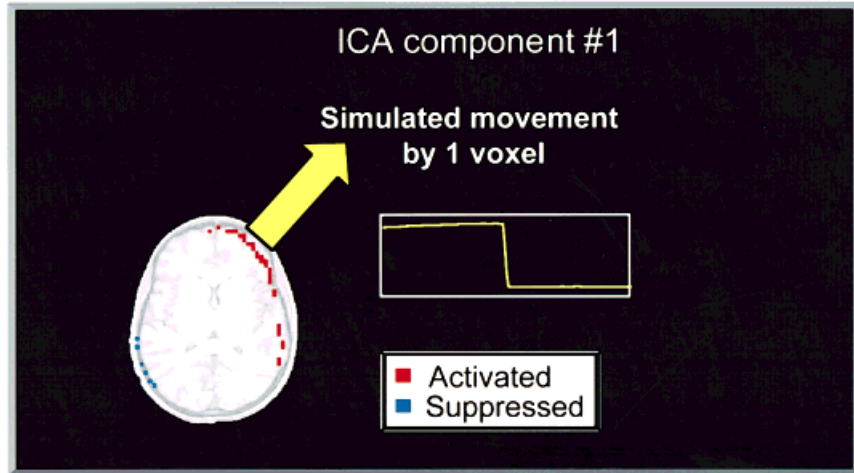
Map Voxel Values

Figure 19.

Stability of the consistently task-related component. Data from one trial (subject 2, trial 1) were divided into two subsets, one containing the odd-numbered time points and the other the even-numbered time points. ICA was performed separately on the two data subsets and compared with results of ICA decomposition of the whole data set (upper left). Each of the two data-subset analyses returned one component with a square-wave time course closely matching that of the CTR component in the analysis of the whole trial (bottom and right). Map voxel values for the square-wave component in each subset analysis were highly correlated with each other and with the map values of the CTR component in the whole-data analysis (scatter plots clipped to $-5 < z < 10$).

The ICA algorithm appears to provide a powerful method for exploratory analysis of fMRI data in both clinical and normal subject populations. It makes no assumptions about the hemodynamic activation function, which may vary across time and brain areas [Kwong et al., 1992; Bandettini et al., 1992]. ICA is at least as sensitive as correlation or PCA in finding task-related activations, and can isolate potentially significant phenomena in the data while canceling out artifacts, using only minimal assumptions about the

Figure 20.

spatiotemporal structure of the component signals. In addition, the ICA method may allow straightforward analysis of more complex brain imaging experiments in which unpredictable changes in cognitive activation occur in parallel with changes in arousal or autonomic states for which the exact time courses of activation are also unknown.

ACKNOWLEDGMENTS

Figure 20.

(a) ICA detection of simulated head movement. The figure shows that the largest component found by ICA after an abrupt head movement was simulated halfway through the time course of one trial by shifting the spatial structure of the data diagonally by one voxel (4.2 mm). The time course of this component clearly mirrors the time course of the simulated movement. The active voxels of this component are those having the largest expected signal change during the movement (the contrast of the structural MRI has been adjusted for clarity). (b) Detection by ICA of a simulated task component. **top:** A small simulated task-related signal (left) consisting of three square waves (instead of four) was either added to (red) or subtracted from (blue) the data from one of the Stroop trials (subject 2, trial 1) in arbitrarily selected portions of two of the functional slices. **middle:** Simulated signal variance was only 0.68% of the mean variance of the arbitrarily selected active voxels (upper left). ICA decomposition of the simulated data recovered two components whose time courses (left) resembled *three* square waves and *four* square waves, respectively. Maps of active voxels ($|z| > 2$) for the three-square-wave component accurately identified the locations and polarities of the simulated active areas (right), with only two false-positive outlying voxels (slice two, bottom left). Correlating the simulated data with a three-square-wave reference function and using a standard correlation threshold ($|r| > 0.4$) detected only 10.7% of the simulated active voxels. **bottom:** When the correlation threshold was reduced until the number of active voxels found by both methods was the same ($n = 195$, $|r| > 0.197$), only 67 (34.4%) of the active voxels selected by correlation were in the simulated active areas, while 128 (65.6%) were false positives.

REFERENCES

Amari S, Cichocki A, Yang H (1996): A new learning algorithm for blind signal separation. Adv Neural Information Processing Syst 8:757–763.

Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, Hyde JS (1992): Time course EPI of human brain function during task activation. Magn Reson Med 25:390–397.

Bandettini PA, Jesmanowicz A, Wong EC, Hyde JS (1993): Processing strategies for time-course data sets in functional MRI of the human brain. Magn Reson Med 30:161–173.

Bell AJ, Sejnowski TJ (1995): An information-maximization approach to blind separation and blind deconvolution. Neural Comput 7:1129–1159.

Bench CJ, Frith CD, Grasby PM, Friston KJ, Paulesu E, Frackowiak RS, Dolan RJ (1993): Investigations of the functional anatomy of attention using the Stroop test. Neuropsychologia 31:907–922.

Binder JR (1997): Neuroanatomy of language processing studied with functional MRI. Clin Neurosci 4:87–94.

Biswal B, DeYoe AE, Hyde JS (1996): Reduction of physiological fluctuations in fMRI using digital filters. Magn Reson Med 35:107–13.

Boynton GM, Engel SA, Glover GH, Heeger DJ (1996): Linear systems analysis of functional magnetic resonance imaging in human V1. J Neurosci 16:4207–4221.

Brammer MJ, Wright IC, Woodruff PWR, Williams SCR, Simmons A, Bullmore ET (1997): Wavelet analysis of periodic and non-periodic experimental designs in functional magnetic resonance imaging of the brain [abstract]. Neuroimage 5:479.

Chichibu S, Ohta Y, Chikugo T, Suzuki T (1995): Temporal and spatial properties of slow waves in the electroencephalogram of spontaneously hypertensive rats. Clin Exp Pharmacol Physiol [Suppl] 22:288–289.

Comon P (1994): Independent component analysis: A new concept? Signal Processing 36:11–20.

Cox RW (1996): AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res 29:162–173.

Duffield JS, de Silva RN, Grant R (1994): Pure alexia without agraphia: A classical cortical syndrome revisited. Scott Med J 39:178–179.

Friston KJ (1995): Commentary and opinion: II. Statistical parametric mapping: Ontology and current issues. J Cereb Blood Flow Metab 15:361–370.

Friston KJ (1996): Statistical Parametric Mapping and Other Analyses of Functional Imaging Data. In: Toga AW, Mazziotta JC (eds): Brain Mapping: The Methods. San Diego: Academic Press, pp 363–396.

Friston KJ, Williams S, Howard R, Frackowiak RS, Turner R (1996): Movement-related effects in fMRI time-series. Magn Reson Med 35:346–355.

Gardner E (1975): Fundamentals of Neurology, 6th ed. Philadelphia: W.B. Saunders.

Jackson JE (1991): A User's Guide to Principal Components, New York: John Wiley & Sons, Inc.

Jung T-P, Humphries C, Lee T-W, Makeig S, McKeown MJ, Iragui V, Sejnowski T (1998): Extended ICA removes artifacts from Electro-encepholographic Recordings. Adv Neural Information Processing Systems, Vol 10 (in press).

Jutten C, Herault J (1991): Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture. Signal Process 24:1–10.

Kwong KK (1995): Functional magnetic resonance imaging with echo planar imaging. Magn Reson Q 11:1–20.

Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, Cheug H-M, Brady TJ, Rosen BR (1992): Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. Proc Natl Acad Sci USA 89:5675–5679.

Le TH, Hu X (1996): Retrospective estimation and correction of physiological artifacts in fMRI by direct extraction of physiological activity from MR data. Magn Reson Med 35:290–298.

Lezak MD (1995): Neuropsychological Assessment, 3rd ed. New York: Oxford University Press.

Makeig S, Bell AJ, Jung T-P, Ghahremani D, Sejnowski TJ (1997): Blind separation of auditory event related responses into independent components. Proc Natl Acad Sci USA 94:10979–10984.

Manoach DS, Schlaug G, Siewert B, Darby DG, Bly BM, Benfield A, Edelman RR, Warach S (1997): Prefrontal cortex fMRI signal changes are correlated with working memory load. Neuroreport 8:545–549.

McKeown MJ, Jung T-P, Makeig S, Brown GG, Kindermann SS, Lee T-W, Sejnowski TJ (1998): Spatially independent activity patterns in functional magnetic resonance imaging data during the Stroop color-naming task. Proc Natl Acad Sci USA 95:803–810.

Mitra PP, Ogawa S, Hu X, Ugurbil K (1997): The nature of spatiotemporal changes in cerebral hemodynamics as manifested in functional magnetic resonance imaging. Magn Reson Med 37:511–518.

Moeller JR, Strother SC (1991): A regional covariance approach to the analysis of functional patterns in position emission tomographic data. J Cereb Blood Flow Metab 11:A121–A135.

Nathaniel-James DA, Fletcher P, Frith CD (1997): The functional anatomy of verbal initiation and suppression using the Hayling test. Neuropsychologia 35:559–566.

Nobre AC, Sebestyen GN, Gitelman DR, Mesulam MM, Frackowiak RS, Frith CD (1997): Functional localization of the system for visuospatial attention using positron emission tomography. Brain 120:515–533.

Ogawa S, Tank DW, Menon R, Ellermann JM, Kim SG, Merkle H, Ugurbil K (1992): Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. Proc Natl Acad Sci USA 89:5951–5955.

Peterson L, Peterson M (1959): Short-term retention of individual verbal items. J Exp Psychol 58:193–198.

Phelps EA, Hyder F, Blamire AM, Shulman RG (1997): FMRI of the prefrontal cortex during overt verbal fluency. Neuroreport 8:561–565.

Phillips CG, Zeki S, Barlow HB (1984): Localization of function in the cerebral cortex. Past, present and future. Brain 107:327–361.

Press WH, Teukolsky SA, Vetterling WT, Flannery BP (1992): Numerical Recipes in C: The Art of Scientific Computing, 2nd ed. Cambridge: Cambridge University Press.

Quint DJ, Gilmore JL (1992): Alexia without agraphia. Neuroradiology 34:210–214.

Stuart A, Ord J (1991): Kendall's Advanced Theory of Statistics, New York: Oxford University Press.

Takahashi N, Kawamura M, Shinotou H, Hirayama K, Kaga K, Shindo M (1992): Pure word deafness due to left hemisphere damage. Cortex 28:295–303.

Tulving E, Markowitsch HJ, Craik FE, Habib R, Houle S (1996): Novelty and familiarity activations in PET studies of memory encoding and retrieval. Cereb Cortex 6:71–79.

Wayenberg JL, Hasaerts D, Franco P, Valente F, Massager N (1995): Anterior fontanelle pressure variations during sleep in healthy infants. Sleep 18:223–228.

Wong E, Bandettini P, Hyde J (1992): Echo-planar imaging of the human brain using a three axis local gradient coil. In: Abstracts of the Proceedings of the Eleventh Annual Meeting of the Society of Magnetic Resonance in Medicine, Berkeley (Society of Magnetic Resonance in Medicine).