Spatial independent component analysis of functional MRI timeseries: To what extent do results depend on the algorithm used?

◆ Human Brain Mapping 16:146–157(2002) ◆

Spatial Independent Component Analysis of Functional MRI Time-Series: To What Extent Do Results Depend on the Algorithm Used?

Fabrizio Esposito,[1]* Elia Formisano,[2] Erich Seifritz,[3] Rainer Goebel,[2] Renato Morrone,[4] Gioacchino Tedeschi,[1] and Francesco Di Salle[5]

[1]*Institute of Neurological Sciences, Second University of Naples, Italy*
[2]*Department of Cognitive Neuroscience, Maastricht University, The Netherlands*
[3]*Department of Psychiatry, University of Basel, Switzerland*
[4]*"Morrone" Diagnostic Center, Caserta, Italy*
[5]*Department of Neurological Sciences, University of Naples "Federico II", Italy*

Abstract: Independent component analysis (ICA) has been successfully employed to decompose functional MRI (fMRI) time-series into sets of activation maps and associated time-courses. Several ICA algorithms have been proposed in the neural network literature. Applied to fMRI, these algorithms might lead to different spatial or temporal readouts of brain activation. We compared the two ICA algorithms that have been used so far for spatial ICA (sICA) of fMRI time-series: the Infomax (Bell and Sejnowski [1995]: Neural Comput 7:1004–1034) and the Fixed-Point (Hyvärinen [1999]: Adv Neural Inf Proc Syst 10:273–279) algorithms. We evaluated the Infomax- and Fixed Point-based sICA decompositions of simulated motor, and real motor and visual activation fMRI time-series using an ensemble of measures. Log-likelihood (McKeown et al. [1998]: Hum Brain Mapp 6:160–188) was used as a measure of how significantly the estimated independent sources fit the statistical structure of the data; receiver operating characteristics (ROC) and linear correlation analyses were used to evaluate the algorithms' accuracy of estimating the spatial layout and the temporal dynamics of simulated and real activations; cluster sizing calculations and an estimation of a residual gaussian noise term within the components were used to examine the anatomic structure of ICA components and for the assessment of noise reduction capabilities. Whereas both algorithms produced highly accurate results, the Fixed-Point outperformed the Infomax in terms of spatial and temporal accuracy as long as inferential statistics were employed as benchmarks. Conversely, the Infomax sICA was superior in terms of global estimation of the ICA model and noise reduction capabilities. Because of its adaptive nature, the Infomax approach appears to be better suited to investigate activation phenomena that are not predictable or adequately modelled by inferential techniques. *Hum. Brain Mapping 16:146–157, 2002.* © 2002 Wiley-Liss, Inc.

Key words: functional magnetic resonance imaging; exploratory data-driven analysis; independent component analysis; Fixed-Point algorithm; adaptive algorithm; information maximization; receiver operating characteristics; Infomax algorithm

INTRODUCTION

Hypothesis-driven and data-driven statistics represent two fundamentally different conceptions in reading out blood oxygenation level dependent (BOLD) functional magnetic resonance imaging (fMRI) time-series. Univariate hypothesis-driven methods [Bandettini et al., 1993; Friston, 1996] provide simple and computationally efficient approaches to produce maps of task-related activations with estimates of their levels of significance. Usually in these methods, an a

Grant sponsor: Swiss National Science Foundation; Grant number: 63-58040.99.

*Correspondence to: Fabrizio Esposito, Division of Neuroradiology, University of Naples "Federico II", II Policlinico (Nuovo Policlinico), Padiglione 16, Via S. Pansini 5, 80127 Naples, Italy.
E-mail: mnlfes@tin.it

Received for publication 16 October 2001; accepted 30 January 2002
DOI 10.1002/hbm.10034

priori temporal model is specified and a statistical parameter is computed at each voxel individually. As a consequence, the specificity of these methods across all the brain regions under investigation is crucially dependent on how correctly such hypotheses are formulated. Additionally, their univariate nature may easily produce a loss of sensitivity if the fMRI experiment induces co-activation of spatially disparate areas with slightly different temporal behaviors.

In contrast, multivariate data-driven techniques enable an exploratory analysis of fMRI data-sets [Lange, 1999] and may potentially separate meaningful activation by computing suitable statistical models independent of any reference paradigm. Furthermore, their multivariate nature exploits the relationships between voxels and may possibly provide useful information about co-activation in spatially different areas of the brain.

Among the data-driven techniques, independent component analysis (ICA) has been shown to provide a powerful method for the exploratory analysis of fMRI data [Arfanakis et al., 2000; McKeown et al., 1998; Moritz et al., 2000]. ICA is an information-theoretic approach of transforming multidimensional data into components that are as statistically independent from each other as possible [Comon, 1994]. As a method for recovering underlying signals, or independent components (ICs) from linear data mixtures, it can be applied to fMRI time-series to spatially localize and temporally characterize the sources of BOLD activation. It is possible to pursue either the temporal [Biswal and Ulmer, 1999] or the spatial [McKeown et al., 1998] independence of the target components. When and how it is useful to apply temporal ICA (tICA) or spatial ICA (sICA) is discussed elsewhere [Calhoun et al., 2001].

Different methods for performing ICA decompositions have been proposed [Giannakopoulos et al., 1999], which use different objective functions (see Materials and Methods) together with different criteria of optimization of these functions, and may potentially produce different results.

The purpose of this study is to evaluate and compare two different neural algorithms for estimating the sICA model on fMRI data: the information maximization (Infomax) approach proposed by Bell and Sejnowski [1995] and a Fixed-Point approach proposed by Hyvärinen [1999]. To this end, we utilized real data from block-designed fMRI activation experiments and simulated signals superimposed on null data. It remains to be elucidated whether a best algorithmic solution for specific applications could be extended to the generality of fMRI time-series [Brown et al., 2001]. Our aim was to explore the performance of different sICA methods on specific experimental designs to provide quantitative information concerning their reliability in the processing of fMRI data, to find out the rules determining differential behaviors of the sICA algorithms, and to propose their differential applications on specific problems.

MATERIALS AND METHODS

Models and algorithms of sICA

Assume $\mathbf{X}$ is a $T \times M$ matrix of observed voxel time courses (T = number of scans, M = number of voxels included in the analysis), $\mathbf{C}$ is a $N \times M$ random matrix whose rows $\mathbf{C}_i$ are to be filled with the (unknown) realizations of the spatial components (N = number of components), and $\mathbf{A}$ is a $T \times N$ "mixing" matrix, whose columns contain the associated time-courses of the N components. The sICA problem for fMRI time-series can be formulated as an estimation of the following linear generative model for the data:

$$X = AC \tag{1}$$

with no assumptions about the mixing matrix $\mathbf{A}$ and with the constraint that the spatial processes $\mathbf{C}_i$ are (ideally) mutually statistically independent (i.e., the joint distribution of an arbitrary group of them factorises [Papoulis, 1991]):

$$p(C_1, C_2, \ldots, C_N) = \prod_{i=1}^{N} p_i(C_i).$$

This is a much stronger criterion than assuming that the $\mathbf{C}_i$ are merely uncorrelated (as in principal component analysis or PCA):

$$C_i C_j^T = \sum_{k=1}^{M} C_{ik} C_{jk} = 0 \quad \text{for } i \neq j.$$

PCA only aims at finding orthogonal spatial maps, or eigenimages, by measuring the "covariance" between all pairs of voxels. It provides second order decompositions through which the data-matrix is summarized by N = T eigenimages, each capturing a certain amount of variance. PCA is commonly used as a dimension reduction technique, in which data are projected onto a reduced set of eigenimages. Whether a dimension reduction is performed or not, PCA is usu-

ally employed as a first step toward statistical independence within a process called "whitening."

The amount of statistical dependence within a fixed number of spatial components can be quantified by means of their mutual information, an important information-theoretic function [Comon, 1994]. Thus, the ICA decomposition of **X** can be defined (up to a multiplicative constant and to the sign) as an invertible transformation:

$$C = W \cdot X \quad (2)$$

where the matrix **W** (also called the unmixing matrix) is determined such that the mutual information of the target components $\mathbf{C}_i$ is minimized (i.e., $\mathbf{C}_i$ are made "as independent as possible"). Matrix **A** is the pseudo-inverse of **W**.

In the first application of ICA to fMRI time-series [McKeown et al., 1998] the Infomax approach was used: the minimization of mutual information was achieved according to the Infomax principle [Bell and Sejnowski, 1995], by maximizing in an adaptive manner (using the natural gradient concept [Amari et al., 1996]) the output entropy of a neural network with as many outputs as the number of ICs to be estimated. This was achieved by using non-linear outputs of the form $g(\mathbf{WX_j})$, $\mathbf{X_j}$ being the generic observed voxel time-course (columns of **X**) and g a non-linear scalar function with suitable properties. The choice of a non-linearity is related to application-specific requirements and depends on the expected distributions of target components. In the context of fMRI, small activity foci in a large volume are usually expected. Therefore the distribution of the target components is assumed to be super-gaussian (sparse). A sigmoidal function $g(.)$ serves this purpose.

The Fixed-Point algorithm [Hyvärinen, 1999] pursues the same goal (i.e., the minimization of mutual information) using the concept of normalized differential entropy or negentropy [Comon, 1994], another information theoretic function, usefully interpreted as a measure of non-gaussianity of a distribution. By expressing the mutual information in terms of negentropy, it is straightforward to demonstrate [Comon, 1994] that finding an invertible transformation **W** that minimizes the mutual information among $\mathbf{C}_i$ is roughly equivalent to finding directions $\mathbf{w_i}$ (rows of **W** in equation (2)) along which the negentropy of the projected data ($\mathbf{w_i}$ **X**) is maximized.

It is possible to use negentropy either as a so-called one-unit objective function and to estimate the ICs one by one (hierarchical approach) or as a multi-unit objective function, estimating all the ICs at the same time (symmetric approach) [Hyvärinen, 1999]. The Fixed-Point algorithm, in both its symmetric and hierarchical version, practically uses simple estimates of negentropy based on the maximum entropy principle [Hyvärinen, 1998]. These approximations result in the use of suitable non-linearities that are practically equivalent to those required in the Infomax implementation when supergaussianity is required. As the name implies, this algorithm is based on Fixed-Point iterations. This means that, instead of using each data observation (i.e., in sICA, each voxel time-course) or small batches of them for updating immediately the estimate of the unmixing matrix **W** (as is the case for the adaptive Infomax algorithm), it uses sample averages computed over a large number of observations (or even all the observations available).

Besides the use of different objective functions and related approximations, the practical difference between the two algorithms consists in the complementary way of updating the unmixing matrix (and, thus, the IC estimates) within each iteration. The updating rule is adaptive in the Infomax approach and non-adaptive in the Fixed-Point approach. In many real world applications, adaptation may have some advantages in accounting for some non-stationary effects during the collection of observations, but it requires the choice of a learning rate sequence (i.e., a step size in the updating formula) that may crucially affect convergence speed. This choice is not needed in a Fixed-Point approach that, consequently, may prove faster and more reliable in achieving convergence [Hyvärinen, 1999].

Image acquisition

Five right-handed healthy volunteers (3 males, 2 females; age 25–40 years) participated in five sessions of a dominant hand finger-tapping fMRI experiment. Images were acquired on a 1.5 Tesla super-conducting SIGNA MR scanner (General Electric Medical Systems, Milwaukee, WI) using a standard circularly polarized head coil. After the acquisition of T_1-weighted structural volumes serving as anatomical reference, 10 contiguous functional slices, positioned parallel to the bicommissural plane and covering the primary motor and supplementary motor areas were acquired using a conventional gradient-echo echo-planar imaging sequence (repetition time 3 sec; echo time 66 msec; delay time 2,000 msec; flip angle 90°, field of view 210*210 mm, matrix 128 × 128, slice thickness 5 mm). The experimental paradigm consisted of ten blocks of five volumes during which a self-paced finger-tapping

task (sequential opposition of all fingers of the right hand against the thumb) at a specified frequency of 2 Hz, was carried out, and ten blocks of five volumes during resting. The frequency of execution was externally controlled by visual inspection.

Additional fMRI data-sets were collected from the five subjects during a visual stimulation, with acquisition parameters and experimental paradigm equivalent to those designed for the motor experiments (ten blocks of 8 Hz photic stimulation and ten blocks of rest, each block lasting five volumes). The ten functional slices were positioned along the bicommissural plane and covered the occipital lobe.

Simulation data

Simulated fMRI time-series were created by adding activation patterns to noise-only datasets (100 volumes) collected in the same subjects during rest (null condition). This allowed taking into account realistic properties of the background noise and avoiding any underlying assumption on its distribution or power. For each subject, the spatial layouts of the simulated activations were obtained from the spatial layouts of the regions activated during the experiment with motor task (primary motor cortex and supplementary motor areas) and detected by conventional linear correlation analysis [Fadili et al., 2000].

At each selected voxel within these regions, the same simulated activation time-course was injected in the null data set using an additive model. For each subject different artificial data-sets were generated; each data set had a different value for the contrast-to-noise ratio (CNR) of the response (CNR = $\Delta S/\sigma_n$, where ΔS is the signal enhancement and σ_n is the noise standard deviation (SD) in those voxels) [Baumgartner et al., 2000].

For the simulation presented thereafter, the selected CNR values were 0.75, 1, 2, and 3 that include the typical CNR range of BOLD contrast at 1.5 T at a conventional voxel size [Bandettini et al., 1992; Kwong, 1995]. The simulated fMRI responses consisted of a boxcar-shaped waveform convolved with a Poisson kernel having parameter $\lambda = 6$ sec [Friston et al., 1994].

Data analysis

Voxels outside the brain were excluded from subsequent analysis by using a histogram-based technique. The remaining voxel-time-courses were used to fill the data matrix **X** in equation (1). To avoid any differential influence of the pre-processing methods on the two ICA approaches, no pre-processing of the data was performed before ICA. The Infomax approach was applied with the same learning rate sequence proposed in [McKeown et al., 1998] and the Fixed-Point algorithm was applied in symmetric mode. A pre-processing step, which included a correction for timing differences between slices, a motion correction, and a temporal high-pass filtering (drift removal), was performed only before the linear correlation analysis of the real data.

The entire data analysis was implemented in Matlab™, using the code for ICA downloaded from the Internet (http://www.cnl.salk.edu for the Infomax approach and http://www.cis.hut.fi for the Fixed-Point approach) and additional code developed in our lab for the purpose of simulations and performance quantification.

For each individual data-set, the matrices **A** and **C** in equation (1) (i.e., the basic model parameters) were repeatedly estimated for 50 runs of the algorithms to cover the effects of random initial conditions for the matrix **W** in equation (2) and then averaged across all runs. The resulting IC values across all the voxels (ICA maps) were, then, scaled to z-scores (i.e., the number of SD from the map mean) as in McKeown et al. [1998].

To assign a statistical significance to these parameters, a gaussian mixture model (GMM) [Beckmann et al., 2001] was adopted and a general gaussian distribution was fitted (using an expectation maximization algorithm) to the histograms of the IC values. This model enables the statistical characterization of a residual (normal-distributed) error term within a single IC: the original IC maps could then be further transformed in new z-maps according to:

$$z_{GMM} = \frac{z - \hat{m}}{\hat{\sigma}} \tag{3}$$

where $\hat{m}$ *e* $\hat{\sigma}$ are the mean and the standard deviation of the fitted normal distribution.

The mixture models have already been used in the context of fMRI activation detection [Everitt and Bullmore, 1999]. Here, this approach is applied only to characterize the null condition ("no activation") within each single IC map and get an estimate of the (marginal) residual noise of the decomposition. This fit can be robust and the estimates of the residual gaussian terms unbiased, only if relatively few voxels are activated (i.e., the IC map is sufficiently "sparse" or "super-gaussian").

Strategies for the quantitation and the comparison of the algorithms' performance

Model estimation: likelihood analysis

As the basic ICA model is a generative model (i.e., there is no error term or any kind of residuals), the "goodness-to-fit" of each decomposition to our data has been evaluated through a likelihood analysis [McKeown and Sejnowski, 1998]. This analysis provides an estimate of how significantly a fixed number of ICs can theoretically predict new observations in the same conditions, assuming that the basic ICA model actually holds.

Formally, we estimated the unmixing matrix **W** of equation (2) by using each of the two ICA algorithms and then we estimated the probability of observing the ith voxel-time-course under the model specified in equation (1), as described in McKeown and Sejnowski [1998]. The minus log-likelihood of the data, computed for each voxel included in the analysis was then averaged across all voxels to provide a global figure of merit for each estimated decomposition [McKeown and Sejnowski, 1998]. This parameter can be computed for an ICA decomposition for a varying number of target components, provided that, for $N < T$, the remaining T–N rows of **W** are filled using the partial PCA decomposition previously removed in a dimension reduction pre-processing step. The case $N = 0$ (i.e., no ICA performed) corresponds to a simple PCA decomposition. Higher values for this figure indicate a lower goodness-to-fit and so a worse model estimation and vice versa. In [McKeown and Sejnowski, 1998], this same likelihood analysis was used to make inferences about the optimal number of principal components to be selected by means of a preliminary PCA. Here, it is used only with the purpose to compare the performances of the two algorithms with a varying number of target components and to evaluate the net gain in likelihood resulting from switching to the higher order statistics (i.e., from PCA–ICA) by using an Infomax or Fixed-Point approach.

Estimation of task-related effects

We evaluated the effectiveness of the two sICA algorithms in estimating the temporal dynamics and the spatial layout of the task related effects using the following methods: the sets of estimated ICs were re-ordered on the basis of the correlation coefficient between the representative time-courses and a suitable reference waveform. Then, the best-ranked IC was selected and labeled as consistently task-related (CTR) component [McKeown et al., 1998]. The value of the cross-correlation coefficient corresponding to this component was used as a temporal figure of merit for the ICA decomposition [Baumgartner et al., 2000].

Spatial accuracy has been evaluated using receiver operating characteristics (ROC) methodology [Skudlarski et al., 1999; Sorenson and Wang, 1996] for either simulated (where we can separate with infinite precision *false* and *true* activation areas) and real activation data. In this latter case, standard linear correlation maps were used as benchmarks.

For all ICA model estimations, an ROC curve has been fitted after the determination of the corresponding false positive fraction (FPF) and the false negative fraction (FNF) at varying thresholds for the selected ICA z-map. The rating scale for ROC calculations has been optimized to provide a good sampling of an acceptable interval within an fMRI working regime (i.e., FPF ranging from 0–0.01) [Fadili et al., 2000] and an ROC power (i.e., the mean of the ROC curve over this range of FPFs) has been utilized as a spatial figure of merit [Skudlarski et al., 1999].

A repeated ROC analysis has been performed on simulated activation data-sets for different values of superimposed CNR and the resulting ROC powers have been averaged across subjects and sessions. The same was done for the corresponding representative time-courses with respect to the computed correlation coefficient. The artificial signal used for generating the simulated activation was adopted as reference.

For all real activation time-series, linear correlation maps with *P*-value thresholds between 0.01–0.00001 were employed as intra-individual benchmarks and used for the definition of true positives and true negatives in the statistical assessment of the two ICA algorithms. More than one threshold was used to cover possible inter-subject variability of responsiveness and thus CNR.

To exploit the filtering effects provided by a multimodal decomposition such as ICA, we checked for distributional and structural properties of ICA maps. If the hypothesized GMM model actually holds, the variances of the estimated gaussian terms of the distributions become a measure of the residual noise after the decomposition.

The spatial structure of the same distributions has been worked out by means of a cluster sizing function: the number of clusters of at least n active voxels ($P = 0.01$ and $P = 0.001$) was plotted against n (minimum cluster dimension) (with n ranging from 2–20 voxels, corresponding to an in-plane extension ranging from 4.88–48.76 mm^2 or to a voxel size ranging from 24.34–243.36 mm^3).

RESULTS AND DISCUSSION

Estimation of the ICA model

The likelihood analysis was performed on real activation data and the results were similar across subjects. For each data set and each run, two main results were observed (Fig. 1). The increase of the number of the principal components selected by the preliminary PCA and used in the ICA decomposition produced a progressive decline of the averaged minus-log-likelihood estimate (corresponding to an improved likelihood of the decomposition) (Fig. 1a). This reduction was always greater if we performed the model estimation by the Infomax algorithm compared to the Fixed-Point algorithm (Fig. 1a,b).

Because the intrinsic dimensionality of fMRI data can affect ICA decompositions, and to avoid any confusion resulting from the filtering effects of partial PCA decomposition, the subsequent ROC investigations were limited for both algorithms to the case of a full ICA decomposition (i.e., no dimension reduction operated by a preliminary PCA).

Characterization of task-related effects

For all subjects, sessions and runs, unique task-related activation maps and associated time-courses of good technical quality were obtained by both algorithms.

Figure 2 shows comparison results of the performances of the two algorithms on the artificial data at four different CNRs of the simulated activation signal (CNR = 0.75, 1, 2, and 3). The measures of spatial and temporal accuracy reveal that the Fixed-Point solution ensured slightly superior performances in terms of ROC power and correlation coefficient; however, the difference between the performances becomes significant only in the case of the highest CNR, suggesting that only strong activation phenomena, less corrupted by noise, tend to be better detected by the Fixed-Point ICA. Figure 2 also shows the absolute ROC power for standard linear correlation maps. It can be observed that correlation analysis always performed better than ICA, in terms of ROC power. This is not surprising because the waveform used for the generation of the artificial data was also used as an input to the linear correlation analysis.

Figure 3 shows ROC results on real motor and visual activation data after averaging across subjects. The indications of better performances of the Fixed-Point algorithm, as suggested by the observed values for the ROC power at different thresholds used to define the "true activation" areas, are quite similar to those observed for the artificial data. The choice of the significance for the linear correlation maps used as benchmarks proved to be a critical parameter for the estimated ROC power.

Because multi-modal data decompositions like ICA present inherent filtering properties not available in standard linear correlation analysis, the ICA and the

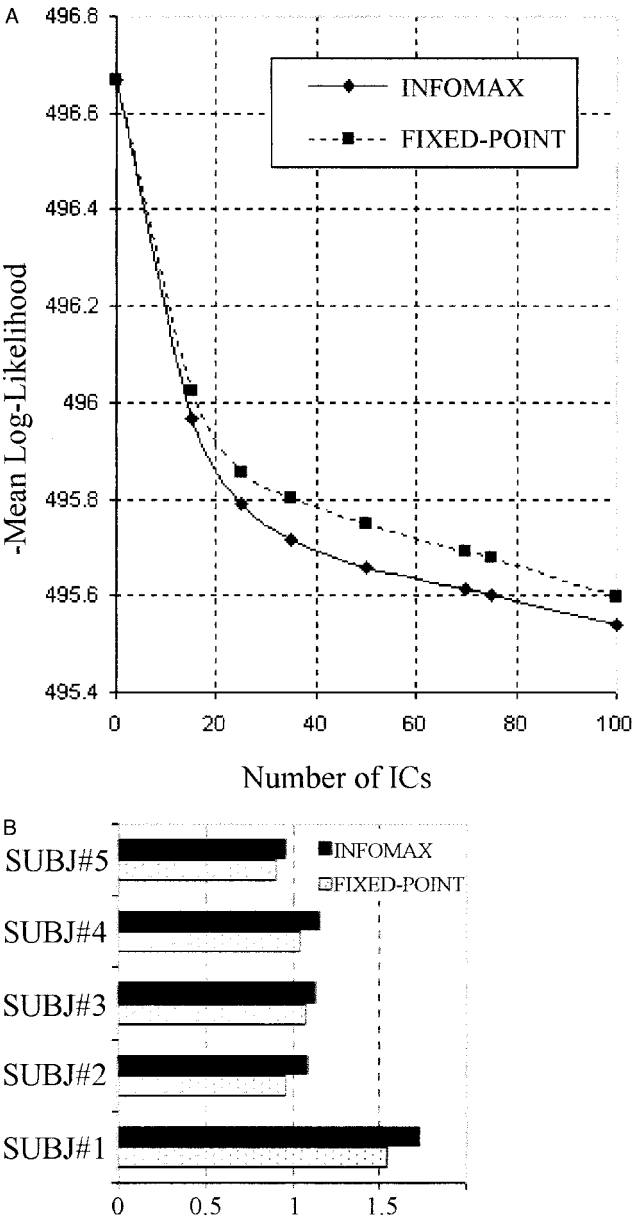

Figure 1.
Results of the likelihood analysis for real motor activation data. **a:** Intra-individual trend of the mean likelihood function for a varying number of dimensions preserved in the pre-whitening step. **b:** Net gain in the mean likelihood for all case subjects.

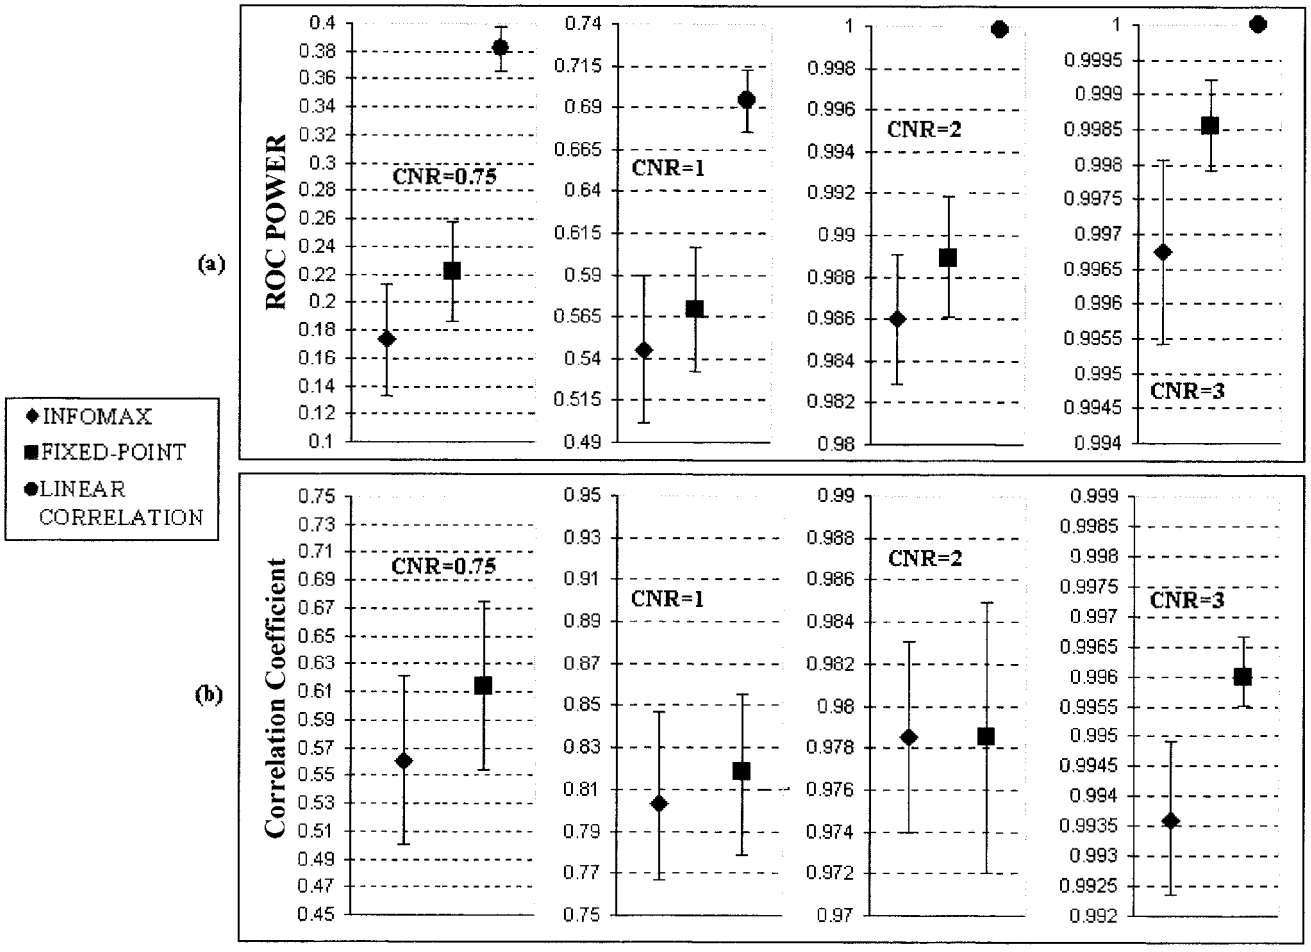

Figure 2.

Results of the comparison between the Infomax and the Fixed-Point ICA algorithms on simulated activation data. **a:** Spatial accuracy (as assessed by ROC analysis) of ICA maps and linear correlation maps. **b:** Temporal accuracy (as assessed by correlation analysis) of the time-courses of the task-related ICA components.

linear correlation analysis exhibit clearly different effects induced by the changes in the threshold. The lower the threshold was set, the more noisy areas were inevitably assumed as "true activation" areas, producing an artifactual reduction of spatial specificity in the ROC analysis of ICA maps (i.e., a greater false positive fraction). The higher the threshold was chosen, the more low contrast active areas were likely to be excluded from the reference patterns, producing an artifactual reduction of spatial sensitivity in the ROC analysis of ICA maps (i.e., a greater false negative fraction).

This phenomenon is also suggested by the motor activation maps shown in Figure 4 for the purpose of a visual comparison. Sensorimotor areas are consistently identified by standard correlation analysis when the reference function models a hemodynamic response sustained throughout the task. Other areas participated in motor functions but may require a modified reference function [Moritz et al., 2000] to be detected by correlation analysis at the same threshold. In particular, the supplementary motor area seems to be heavily under-represented by the correlation analysis compared to ICA, and high contrast foci in primary motor cortex tend to reduce their spatial extension in the correlation analysis, whereas ICA maps show this "primary" activation with a spatial extension matching more closely the gyral anatomy of the Rolandic region.

The sub-optimality of linear correlation analysis has been previously suggested [Friston, 1998], and we have to remark that here the linear correlation is used

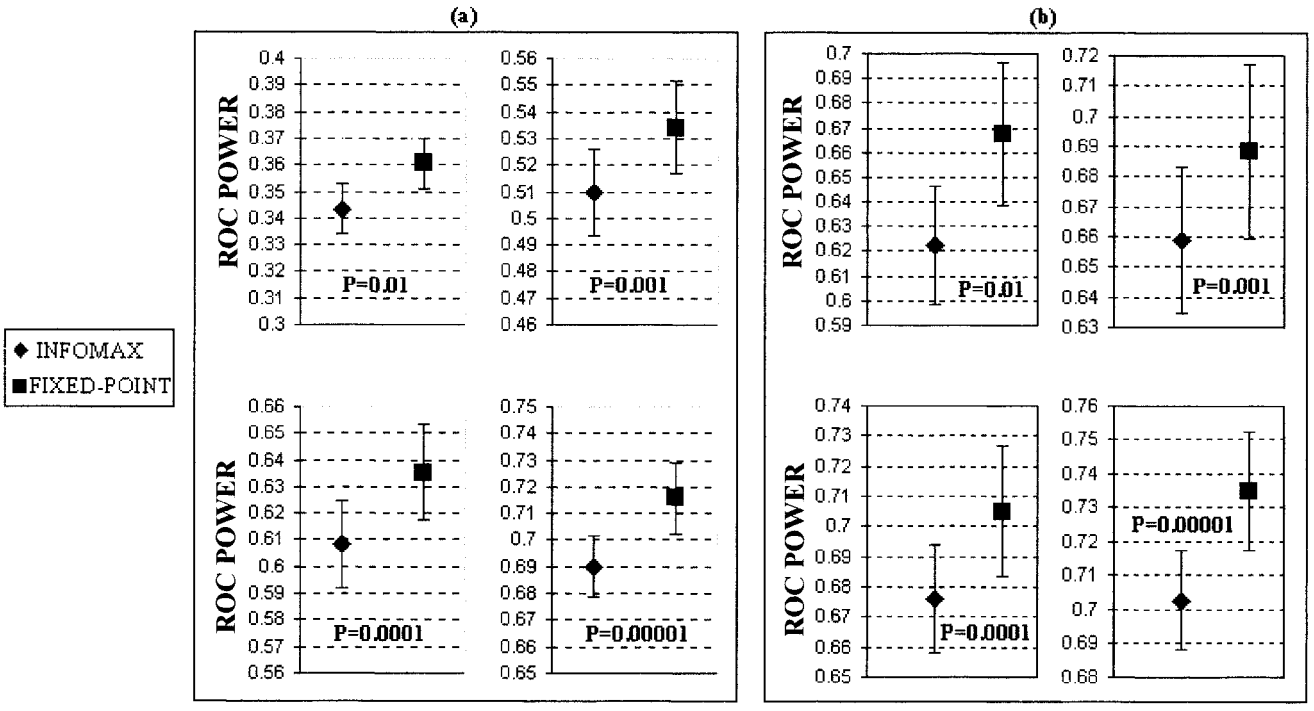

Figure 3.
Results of the comparison between the Infomax and the Fixed-Point ICA algorithms on real activation data. Spatial accuracy of ICA maps as assessed by ROC analysis using correlation map with different thresholds ($P = 0.01–0.00001$) as benchmarks. **a:** Results on motor activation data. **b:** Results on visual activation data.

only in a relative manner as the simplest standard for comparison purposes, to facilitate the interpretation of the results.

Accepting the known sub-optimality of correlation, the results from the ICA analysis of real activation have proven to be quite consistent with simulated activation results in that they show a better spatial coincidence between the standard linear correlation and the Fixed-Point ICA than the Infomax results.

The residual error variances, estimated on ICA maps according to the Gaussian mixture model (without any benchmark), revealed a slightly better performance of Infomax in noise reduction capabilities (i.e., the residual variances of the Fixed-Point components were always superior to those of the Infomax components).

Figure 5 summarizes the effects of gaussian mixture model assumptions on a typical histogram for an ICA component as well as the intra-individual comparison results between Infomax and Fixed-Point components.

It is interesting to observe how, because of the super-gaussianity of the ICA components, their normalized histograms are always more sharply peaked than a standard Gaussian distribution (i.e., mean 0 and SD 1) and, thus, the fitted Gaussian distributions always show a variance <1.

In addition to the greater noise reduction capabilities, Infomax generated more structured components in terms of degree of clustering as well. Figure 6 reports cluster sizing calculations for CTR-ICA maps (generated by both the two approaches) and linear correlation maps, for two values of type I error probability. The filtering effects provided by ICA decompositions are quite evident: even if no structural information is taken into account in the generation of the spatial distribution of the ICA components, the clustering properties of ICA activation maps are always above those of correlation maps.

CONCLUSIONS

We have experimentally compared the two ICA algorithms, Infomax and Fixed-Point, that have been adopted in the fMRI literature. Our purpose was to confirm their robustness and reliability in extracting task-related activation maps and time-courses from single-condition block-designed fMRI data sets. We have chosen a very simple experimental framework to

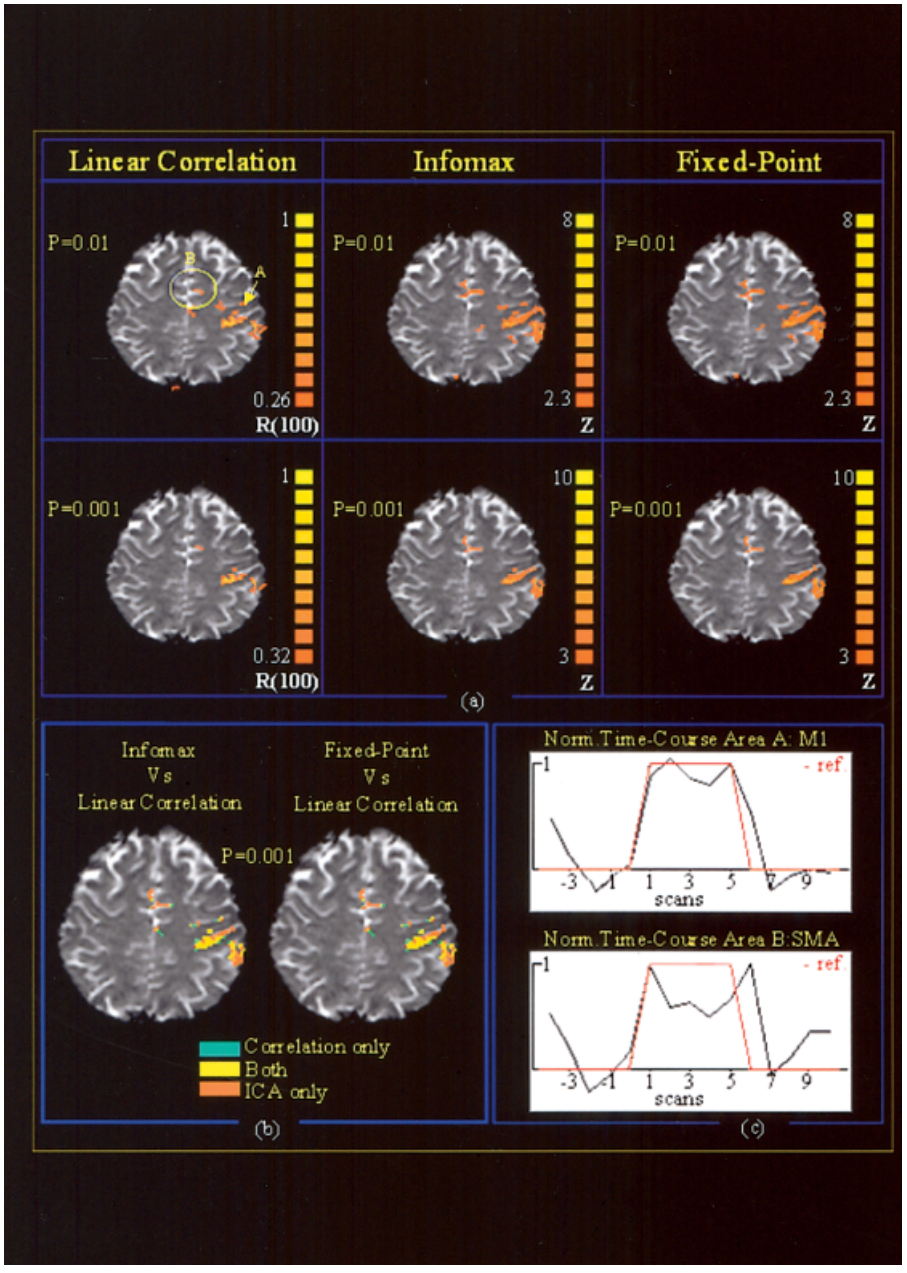

Figure 4.
a: Activation maps from real motor activation data as obtained by linear correlation analysis, Infomax sICA and Fixed-Point sICA (task-related components) at two different thresholds ($P = 0.01$, $P = 0.001$). **b:** Comparison of the results obtained by the two algorithms and linear correlation analysis (P=0.001). **c:** Averaged normalized time-courses from area A, primary motor cortex (M1), and area B, supplementary motor area (SMA) labeled in (a).

cover with great details the most context-specific aspects of using ICA for fMRI time-series. Thereby we intended to investigate algorithms rather than models, for which more challenging experimental designs would be needed. We tried to outline the most refined features of our results and compare distributional and structural properties of ICA decompositions such as those obtained by different algorithmic approaches to the same problem. The results on real motor activation data demonstrate how ICA could yield additional amounts of information even in those situations where a major part of results can be easily expected and easily confirmed by means of standard univariate statistics. This aspect could have very interesting implications in more complex experimental conditions.

The main difficulty with assessing precisely temporal and spatial accuracy on data with real activation responses was concerned with the actual lack of a reliable benchmark: the properties of the true data-generating mechanism, are rarely, if ever, really known, and the experimenter is generally forced to accept some statistical compromise between known and unknown as a pragmatic course of action. We accepted the sub-optimality of linear correlation maps

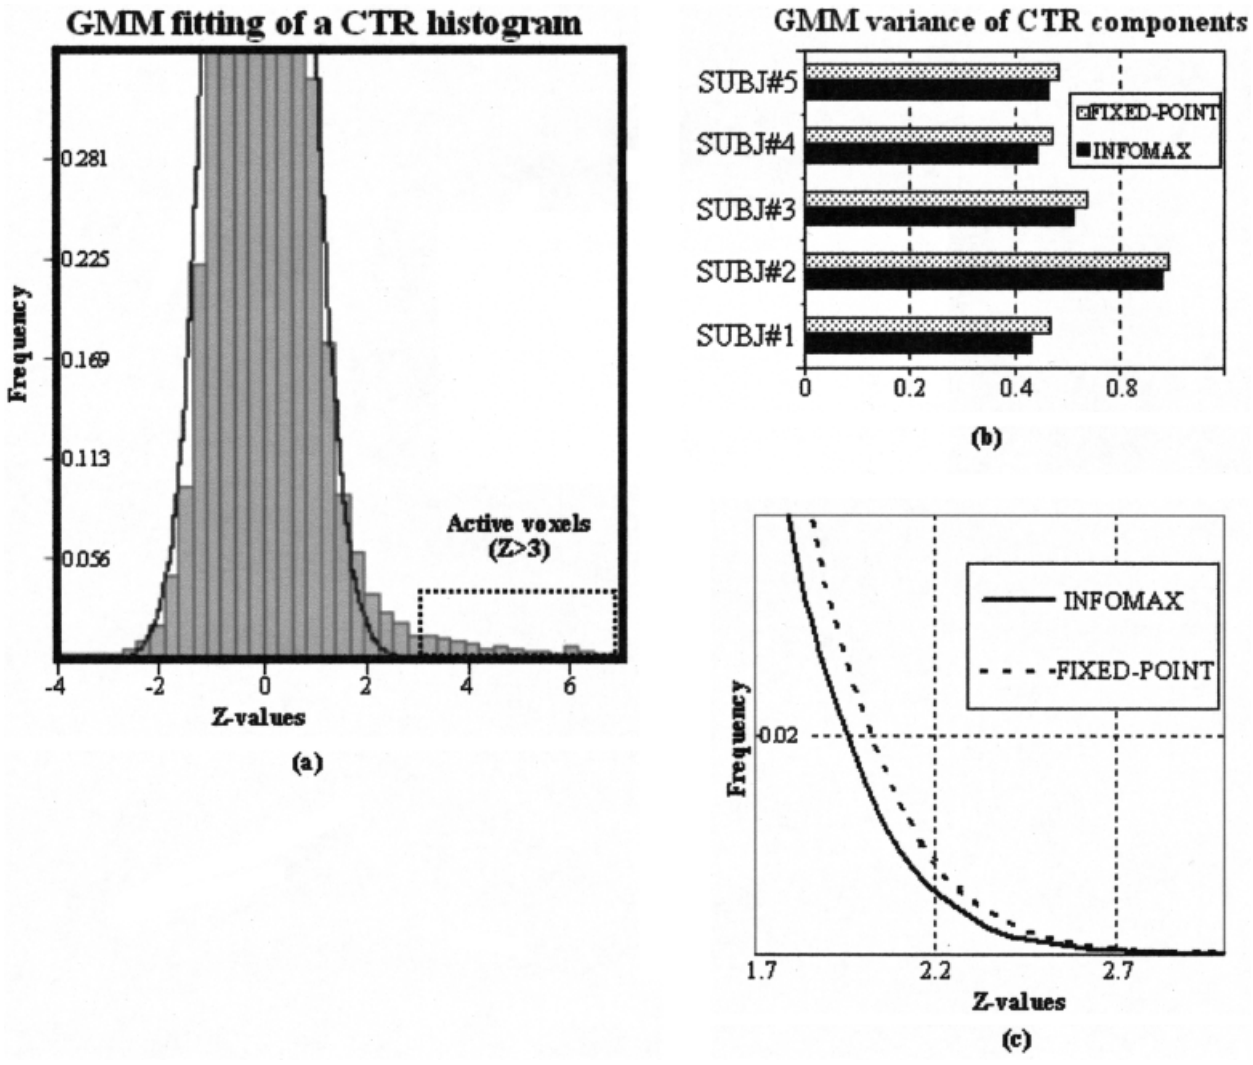

Figure 5.

a: Results from fitting a gaussian distribution to the histogram of a consistent task-related (CTR) sICA component. Note that here the IC is positive signed and active voxels remain outside the area under the fitted gaussian density function. **b:** Variances of the fitted gaussian distribution (residual noise) for the CTR components generated by the two algorithms across all the case subjects (motor activation data). **c:** Comparison of two fitted gaussian distributions.

in our ROC calculations because it is the standard way of treating fMRI data. The estimates of performances we found on real activation data were consistent with those produced on simulated data, indicating that the results of the Fixed-Point approach, compared to the Infomax approach, always produced results that were more overlapping with the correlation analysis results. In contrast, likelihood results showed always better performances for the ICA decompositions generated by Infomax. The global figure of merit provided by likelihood analysis found an interesting confirmation in the GMM characterization of task-related components: after GMM fitting we always found a lower variance for the residual noise in the Infomax generated CTR component with respect to the Fixed-Point approach. If the GMM model actually holds, this means a greater number of potentially "significant" voxels (at fixed Type I error probabilities) generated by Infomax with respect to Fixed-Point, which indeed have been shown to contribute to the anatomical structure of the activation maps by cluster sizing measures.

In conclusion, both algorithms worked properly on our data. The adaptive solution implemented in Info-

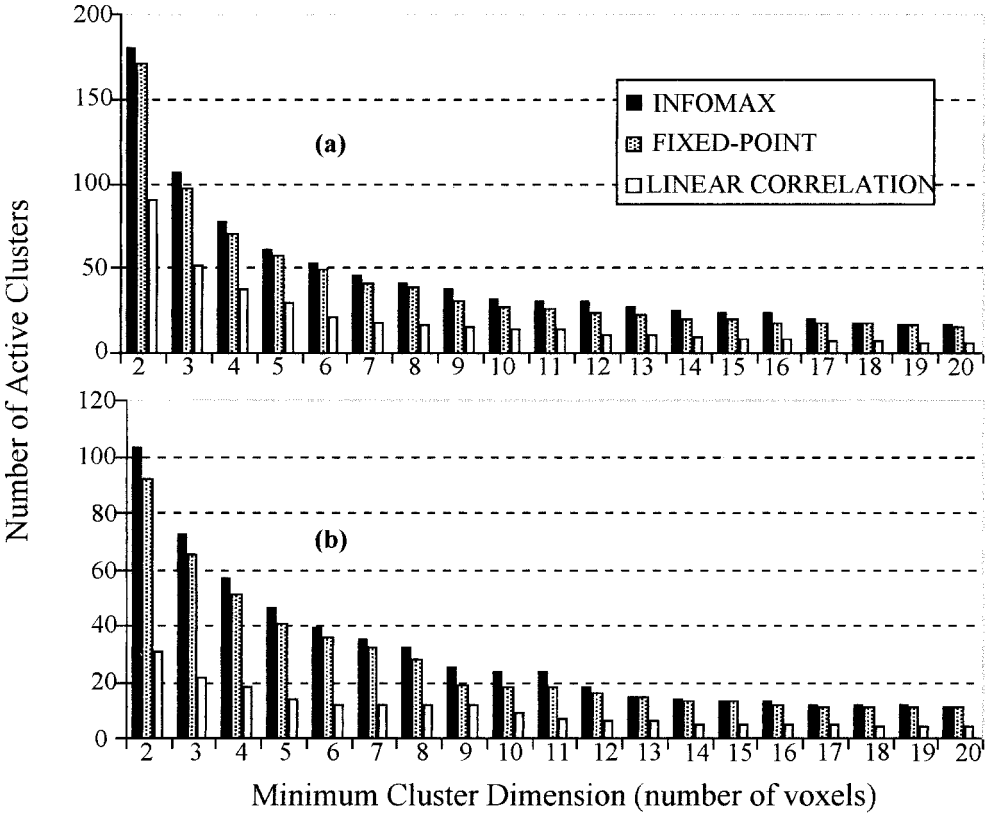

Figure 6.
Cluster sizing measures for the CTR-sICA maps and linear correlation maps. The number of active clusters (**a:** $P = 0.01$; **b:** $P = 0.001$) is plotted against the minimum cluster dimension (in number of voxels).

max showed a slight superiority in global model estimation and filtering capabilities. The better global features of the decompositions of our data provided by Infomax confirm its advantages compared to Fixed-Point in minimizing the mutual information [Giannakopoulos et al., 1999]. The better noise-reduction effects are likely to be related to Infomax's adaptive nature, with the consequence of a higher capability to overcome possible spatial non-stationarities in the fMRI observations.

REFERENCES

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

Arfanakis K, Cordes D, Haughton VM, Moritz CH, Quigley MA, Meyerand ME (2000): Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. Magn Reson Imaging 18:921–930.

Bandettini PA, Jesmanowicz AJ, 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.

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.

Baumgartner R, Ryner L, Richter W, Summers R, Jarmasz M, Somorjai R (2000): Comparison of two exploratory data analysis methods for fMRI: fuzzy clustering vs. principal component analysis. Magn Reson Imaging 18:89–94.

Beckmann CF, Noble JA, Smith SM (2001): Spatio-temporal accuracy of ICA for fMRI. Neuroimage 13:S75.

Bell AJ, Sejnowski TJ (1995): An information–maximization approach to blind separation and blind deconvolution. Neural Comput 7:1004–1034.

Biswal BB, Ulmer JL (1999): Blind source separation of multiple signal sources of fMRI data sets using independent component analysis. J Comput Assist Tomogr 23:265–271.

Brown GD, Yamada S, Sejnowski TJ (2001): Independent component analysis at neural cocktail party. Trends Neurosci 24:54–63.

Calhoun VD, Adali T, Pearlson GD, Pekar JJ (2001): Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. Hum Brain Mapp 13:43–53.

Comon P (1994): Independent component analysis, a new concept? Signal Process 36:287–314.

Everitt BS, Bullmore ET (1999): Mixture model mapping of brain activation in functional magnetic resonance images. Hum Brain Mapp 7:1–14.

Fadili MJ, Ruan S, Bloyet D, Mazoyer B (2000): A multi-step unsupervised fuzzy clustering analysis of fMRI time series. Hum Brain Mapp 10:160–178.

Friston KJ (1996): Statistical parametric mapping and other analyses of functional imaging data. In: Toga AW, Mazziotta JC, editors. Brain mapping: the methods. San Diego: Academic Press. p 363–396.

Friston KJ, Jezzard P, Turner R (1994): Analysis of functional MRI time series. Hum Brain Mapp 1:153–171.

Friston KJ (1998): Modes or models: a critique on independent component analysis for fMRI. Trends Cognit Sci 2:373–375.

Giannakopoulos X, Karhunen J, Oja E (1999): An experimental comparison of neural algorithms for independent component analysis and blind separation. Int J Neural Syst 9:99–114.

Hyvärinen A (1998): New approximations of differential entropy for independent component analysis and projection pursuit. Adv Neural Inf Proc Syst 10:273–279.

Hyvärinen A (1999): Fast and robust Fixed-Point algorithms for independent component analysis. IEEE Trans Neural Networks 10:626–634.

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

Lange N (1999): Statistical procedures for functional MRI. In: Moonen CT, Bandettini PA, editors. Functional MRI. New York: Springer. p 301–335.

McKeown MJ, Makeig S, Brown GG, Jung T-P, Kindermann SS, Bell AJ, Sejnowski TJ (1998): Analysis of fMRI data by blind separation into spatial independent component analysis. Hum Brain Mapp 6:160–188.

McKeown MJ, Sejnowski TJ (1998): Independent component analysis of fMRI data: examining the assumptions. Hum Brain Mapp 6:368–372.

Moritz CH, Haughton VM, Cordes D, Quigley M, Meyerand E (2000): Whole-brain functional MR imaging activation from a finger-tapping task examined with independent component analysis. Am J Neuroradiol 21:1629–1635.

Papoulis A (1991): Probability, random variables, and stochastic processes. New York: McGraw-Hill.

Skudlarski P, Constable RT, Gore JC (1999): ROC analysis of statistical methods used in fMRI: individual subjects. Neuroimage 9:311–329.

Sorenson JA, Wang X (1996): ROC method for evaluation of fMRI Techniques. Magn Reson Med 36:737–744.