

Independent component analysis of fMRI group studies by self-organizing clustering

Fabrizio Esposito,^{a,b,c,*} Tommaso Scarabino,^b Aapo Hyvarinen,^d Johan Himberg,^e
Elia Formisano,^f Silvia Comani,^g Gioacchino Tedeschi,^a Rainer Goebel,^f
Erich Seifritz,^h and Francesco Di Salle^{b,c}

^aSecond Division of Neurology, Second University of Naples, Italy

^bDepartment of Neuroradiology, Scientific Institute “Casa Sollievo della Sofferenza”, S. Giovanni Rotondo, Italy

^cDepartment of Neurological Sciences, University of Naples “Federico II”, Italy

^dHIIT Basic Research Unit, Department of Computer Science, University of Helsinki, Finland

^eNeural Networks Research Centre, Helsinki University of Technology, Finland

^fDepartment of Cognitive Neuroscience, Maastricht University, The Netherlands

^gDepartment of Clinical Sciences and Bio-imaging, University of Chieti, Italy

^hUniversity Hospital of Clinical Psychiatry, University of Bern, Switzerland

Received 14 July 2004; revised 21 October 2004; accepted 28 October 2004
Available online 8 January 2005

Independent component analysis (ICA) is a valuable technique for the multivariate data-driven analysis of functional magnetic resonance imaging (fMRI) data sets. Applications of ICA have been developed mainly for single subject studies, although different solutions for group studies have been proposed. These approaches combine data sets from multiple subjects into a single aggregate data set before ICA estimation and, thus, require some additional assumptions about the separability across subjects of group independent components. Here, we exploit the application of similarity measures and a related visual tool to study the natural self-organizing clustering of many independent components from multiple individual data sets in the subject space. Our proposed framework flexibly enables multiple criteria for the generation of group independent components and their random-effects evaluation. We present real visual activation fMRI data from two experiments, with different spatiotemporal structures, and demonstrate the validity of this framework for a blind extraction and selection of meaningful activity and functional connectivity group patterns. Our approach is either alternative or complementary to the group ICA of aggregated data sets in that it exploits commonalities across multiple subject-specific patterns, while addressing as much as possible of the intersubject variability of the measured responses. This property is

particularly of interest for a blind group and subgroup pattern extraction and selection.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Functional magnetic resonance imaging; fMRI; Exploratory data-driven analysis; Descriptive statistics; Multidimensional projection; Spatial; Temporal similarity analysis; Independent component analysis; Group analysis; Multi-subjects analysis

Introduction

Independent component analysis (ICA) is a valuable tool for the multivariate data-driven analysis of functional magnetic resonance imaging (fMRI) data (McKeown et al., 1998a,b, 2003). As a purely data-driven methodology, ICA does not require the specification of temporal signal profiles or anatomical regions of interest to generate meaningful spatiotemporal patterns of brain activity. The multivariate statistical nature of ICA allows one to transform three-dimensional fMRI data sets into brain activity patterns starting from the spatial or temporal covariance of the measured signals and reveals multiple spatiotemporal “modes” of signal variability (Friston et al., 1993). This transformation is achieved by imposing the general, yet neurophysiologically plausible, constraint of removing the statistical dependence of the output modes (Brown et al., 2001; McKeown et al., 1998a). In order to meet this constraint, the value distribution of the fMRI signals in space or time is to be considered: the variant called spatial ICA (sICA) refers to the statistical distribution of signals across the sampled hemodynamic locations, while the variant called temporal ICA

* Corresponding author. Department of Neurological Sciences, University of Naples Federico II, II Policlinico (Nuovo Policlinico) Padiglione 17, Via S. Pansini 5, 80131 Naples, Italy. Fax: +39 081 546 3663.

E-mail address: faesposi@unina.it (F. Esposito).

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

(tICA) refers to the statistical distribution of source signals across the sampled time-points (Calhoun et al., 2001a).

Both sICA and tICA have been used in different contexts. The tICA is applied to fMRI measurements in the same way, and with the same assumptions that ICA is commonly applied to EEG or MEG recordings; on the other hand, the structure of whole-brain three-dimensional fMRI data sets has suggested the sICA as the default ICA variant for fMRI.

The neurological significance of applying sICA for the decomposition of single-subject fMRI time-series can be seen in the two equivalent formulations of ICA. First, the modes of signal change separated by ICA are such that the mutual information is minimized, that is, each generated pattern carries a minimum amount of information about the other patterns. Second, the statistical distribution of the sources are maximally far from the Gaussian distribution (Hyvarinen et al., 2001). In fact, the first definition extends the concept of functional connectivity patterns of brain imaging data (Friston et al., 1993), where multiple brain regions are unified by their time-courses, with the constraint that none of these regions systematically occurs in two different patterns. The second definition fits with the concept of “activation map”, for which the amount of functional information is related to how the values of a few (active) voxels are significantly different from the remaining (Gaussian-distributed, as default) “mass” of voxels: the more Gaussian the distribution of a three-dimensional map is, the less selective and, thus, functionally uninformative, the resulting pattern will be.

Based on this theoretical background, ICA has been successfully used for the decomposition of individual fMRI time-series. However, since fMRI studies increasingly involve the statistical comparison of more than one group of subjects, for example, healthy people vs. people with a disease, it has become necessary to develop strategies to extending the ICA analysis framework from single-subject to group studies and multi-group studies.

The most natural and intuitive way that avoids additional assumptions for the individual ICA data model is to perform fixed- or random-effects analyses on the results of the decompositions of each individual data sets (Calhoun et al., 2001b; Seifritz et al., 2002). The main challenge of this approach is to integrate the ICA analysis chain with a suitable post-estimation analysis step in which an automatic tool would allow a systematic matching of the estimated components across all the subjects of the study. However, in previous studies applying ICA to fMRI, the matching of the component maps was based on subjective and context-specific criteria: in the absence of general and effective tools for the subject- or group-level selection of “matching” components, this approach remains difficult to implement, and the loss of sensitivity caused by a possible mismatch of components cannot be easily corrected.

Conventional model-driven univariate methods (e.g., regression analysis) have been naturally generalized from single- to multi-subject methods by simple schemes of across-subject data aggregation based on matrix averaging or concatenation. Previous work has proposed similar schemes to combine the individual data sets into a single group data set prior to performing one single ICA run on a group data-matrix. Two alternative approaches have been proposed. Following the typical matrix notation, they can be referred as column-wise (or subject-wise) (Calhoun et al., 2001c), and row-wise (across time-courses) concatenation (Svensen et al., 2002). These methods have been reviewed and compared using artificial data sets to the simplest

across-subject averaging by Schmithorst and Holland (2004). In order to be correctly applied, both approaches require the substantial assumption that a given source of signal change exists as an “observable process” in all of the subjects entering the analysis. Specifically, column-wise aggregation imposes a common space of observations for all the sources (the normalized anatomical space), although it allows different activation time-courses for the different subjects. Row-wise aggregation imposes a common time-course for a generic source to all of the subjects, although it allows “no activity” to occur in some of subjects. Despite the additional, sometimes restrictive hypotheses required by the aggregate approaches, the use of a common space of observation may serve as useful “regularization” for the estimation of group components.

After ICA parameter estimation, the separation of subject-specific components is achieved by a subject-level unmixing of group components in the column-wise approach and by a vector disaggregation of group components in the row-wise approach. In a more recent work (Calhoun et al., 2004), a new variant of the column-wise group ICA approach was presented, where single-subject component time-courses were obtained using a spatial multiple regression of the group component images onto the individual fMRI data for each time point.

A further approach is the simple across-subject averaging (Schmithorst et al., 2004): although the computational load is the least extensive and is independent from the number of subjects, it allows group inferences only indirectly through a subsequent conventional general linear model analysis with the estimated ICA mixing matrix acting as a pseudo design matrix in a way similar to that described by McKeown (2000). In all of the three approaches, at least one form of “non-selective” pooling of different subjects’ data is necessary before estimating group components: spatial for column-wise, temporal for row-wise and spatiotemporal in across-subject averaging.

Although the validity of these approaches in producing “single-group” ICA patterns compatible with individual ICA patterns has been demonstrated, it is noteworthy that they cannot easily predict how much bias or loss of sensitivity may occur in the ICA estimation (and, thus, subsequent random effects analysis of the patterns) in the presence of factors affecting the homogeneity across subjects of the components. Thus, the homogeneity of the sample of subjects, which crucially affects the performances of the random effects analysis for model-driven parametric estimates (Friston et al., 1999), appears to be an even more crucial problem in the context of ICA.

The problem of the homogenous presence of sources in different subjects may occur for many different reasons: for instance, Burbaud et al. (2000) show different activation patterns for mental calculation relating to different strategies (verbal or visual), while in the study of Castelo-Branco et al. (2002), an individual could or could not produce a measurable response related to his/her subjective perception of an ambiguous stimulus.

So far, preliminary attempts to examine the homogeneity (or stationarity) of subject (or timepoint-to-timepoint) homogeneity have been presented by Liao et al. (2004) and Calhoun et al. (2001d).

In general, both predictable (e.g., gender, age, etc.) as well as not easily predictable factors may occur, which can bias the group ICA model estimation, but a comprehensive evaluation of this bias and the possible loss of accuracy of the proposed ICA method is not straightforward.

Recent work has suggested a method to assess the homogeneity of the sample of subjects before general linear model random effects analysis, using similarity measures and multidimensional projection of single subject data sets (Kherif et al., 2003); other recent work has shown how single-subject ICA estimates and intersubject correlation can help to dissect the cerebral cortex into stimulus-driven functional connectivity patterns even in highly complex naturalistic settings (Bartels and Zeki, 2004).

Here, we propose the application of similarity measures on ICA patterns to produce group inferences in multi-subject studies: starting from single-subject ICA runs, we explore the natural self-organizing clustering of components in the subject space, assuming the inter-subject similarity, contrasted to the intra-subject similarity, of the component estimates as a cluster generator. We call this approach “self-organizing group ICA” (sogICA), since it extends ICA from individual to multi-subject fMRI data sets without forcing a specific homology of the sources across subjects. In contrast, sogICA searches for structures of the sources in the subject space.

We present all the steps of this framework and show results obtained from real activation fMRI experiments conducted on a group of six subjects. For these experiments, two illustrative experimental paradigms involving visual stimulation have been adopted in a way that either one single or two spatially non-systematically overlapping (i.e., spatially independent) sources were to be expected from each single-subject decomposition under normal conditions.

Materials and methods

Experimental design

Six healthy volunteers, two males and four females (mean age 26 years) with normal vision and audition, were recruited for the study. The Ethical Committee approved the protocol and the participants signed their informed consent. Two experimental paradigms depicted in Fig. 1 were used for two successive 3 min lasting scanning sessions. According to an ON–OFF blocked design scheme, periods of passive visual stimulation were administered to the volunteers while they were laying supine in the scanner with their eyes fixating a white cross in the centre of their visual field: an LCD projector was used to project the visual scenario onto a screen which covered about 25° of the visual field of participants through a mirror. The visual stimulation consisted of four 20-s lasting periods of a 10-Hz flickering checkerboard, covering the entire stimulated visual field in first session and alternatively left and right part of the stimulated visual field in the second session, were alternated to simple fixation 20-s lasting

periods. The stimuli were timed using the computer program Presentation (www.neurobs.com).

Imaging parameters and pre-processing

MR images were acquired on a 3-T super-conducting SIGNA LX scanner (General Electric Medical Systems, Milwaukee, WI, USA) using a standard circularly polarized head coil. A sagittal localizer was performed first that served to position 14 oblique slices parallel to the AC–PC plane, covering the entire brain including the superior part of the cerebellum. Echo-planar imaging (TR = 1 s, TE = 30 ms, FOV = 240×240 mm², matrix 64×64 , slice thickness 5 mm and gap 1.5 mm) was performed for dynamic scans: time-series were acquired for 200 s, the first 20 scans being “dummy” scans that allowed the longitudinal equilibrium of the magnetization vector; after the dummy scans, the program started the stimulation paradigm, automatically triggered by the scanner. After functional scans, co-registered anatomical slices and whole-head T1-weighted isometric (1 mm³) scans were acquired for all subjects for a correct 3-D registration and isotropic time-series generation.

Functional image time-series were first corrected for the differences in slice acquisition times using a “sinc” interpolation technique, realigned with T₁-volumes, warped into the standard anatomical space of Talairach and Tournoux (1988) and finally resampled to $3 \times 3 \times 3$ mm³ isotropic voxels.

The Talairach transformation was performed in two steps. The first step consisted in rotating the 3-D data set of each subject to be aligned with stereotaxic axes (for this step the location of the anterior commissure [AC], the posterior commissure [PC] and two rotation parameters for midsagittal alignment were specified manually). In the second step, the extreme points of the cerebrum were specified. These points together with the AC and PC coordinates were then used to scale the 3-D data sets into the dimensions of the standard brain of the Talairach and Tournoux (1988) atlas using a piecewise affine and continuous transformation for each of the 12 defined subvolumes.

The resulting voxel-time-series (180 scans) were filtered in time and space: low-frequency (drift) fluctuation were reduced using a high-pass temporal filter (5 cycles) and high-frequency fluctuations were reduced using a 4-s full-width at half-maximum Gaussian kernel; spatial smoothing was performed using a 6-mm full-width at half-maximum Gaussian kernel.

Time and space filtering of image time-series are utilized here, as it is common use in fMRI data analysis, to strengthen signal detection and improve the matching of the regions of activity in multi-subject studies. It is straightforward to demonstrate that any spatial (linear) filtering of the data applied before ICA does not change the spatial (linear) ICA model (Hyvarinen et al., 2001).




Fig. 1. Experimental design.

Nonetheless, the impact of spatial and temporal filtering on ICA source separation and estimation has not been studied analytically in fMRI and may deserve further methodological investigation.

No correction for head motion was performed for reasons similar to those discussed previously (Esposito et al., 2003). Namely, since individual ICA detects gradual and sudden motion without a preliminary correction of the data for the subjects' confined head motion (McKeown et al., 1998b), with the components of interest proving adequately robust to motion effects, we avoided to alter the structure of the data sets (by the typical image resampling step required in the common motion correction procedures) in a way that could be highly different across subjects because of the presumably different motion.

For display on the volumetric anatomy, individual maps were either projected on the respective individual volumetric images or on a standard anatomical Montreal Neurological Institute (MNI) template. Surface projections on the inflated cortical surface obtained from the segmentation and unfolding of the MNI template were used for a 3-D comparison of both individual and group maps.

All the image data preparation and preprocessing steps as well as the map volumetric and surface projection were performed in Brain Voyager QX (Brain Innovation, Maastricht, The Netherlands).

ICA data model and framework in fMRI

ICA decomposes data into signals that are maximally statistically independent (Hyvarinen et al., 2001). In fMRI, given P voxels in the brain and T time-points, the basic formulation of the generative data model of the noise-free spatial ICA assumes that the single-subject data matrix \mathbf{X} of order $(T \times P)$, whose columns are the voxel-time courses, can be represented through a linear combination of M statistically independent components (the “sources”) in the P -dimensional space domain, whose observations are organized as rows of a $M \times P$ matrix \mathbf{C} :

$$\mathbf{X} = \mathbf{A} \mathbf{C} \quad (1)$$

\mathbf{A} is the “mixing” matrix $(T \times M)$. The i th column of \mathbf{A} consists of the time-course associated with the i th row of \mathbf{C} , C_i . Most of ICA algorithms compute the “unmixing” matrix \mathbf{W} so that the formulas:

$$\mathbf{C} = \mathbf{W} \mathbf{X}; \mathbf{A} = (\mathbf{W}^T \mathbf{W})^{-1} \mathbf{W}^T \quad (2)$$

can be used for an estimate of the original component sources and their time-courses of activation. The resulting IC map values of all C_i are conventionally scaled to spatial z scores (i.e., the number of standard deviations from the map mean) as in (McKeown et al., 1998a) before projection onto the anatomical space.

The most popular algorithms for the estimation of matrix \mathbf{W} are the Infomax algorithm, which uses a stochastic gradient-descent adaptive algorithm to minimize the mutual information (Bell and Sejnowski, 1995) of the components and the FastICA algorithm, a fixed-point iterative algorithm that maximizes the negative normalized entropy (negentropy) of the components (Hyvärinen, 1999). These algorithms have been reviewed and compared on artificial and real activation fMRI data by Esposito et al. (2002).

In real world applications, it is typical to include in the core ICA estimation framework a preliminary dimension reduction of the input data-matrix \mathbf{X} : a principal component analysis (PCA)-

based reduction of the matrix \mathbf{X} is commonly performed from order $T \times P$ into order $M \times P$. Although a number of model-selection criteria have been proposed in the signal processing literature, heuristic choices of the number of principal-independent components (M) to be estimated were successful in many previous applications of ICA to fMRI. Heuristic choices handle the fundamental methodological trade-off between convergence problems and separation performances: an exaggerated dimension reduction may generate representations of not completely unmixed sources as well as a loss of interesting sources (“under-decomposition”); on the other hand, the inclusion of a number of dimensions higher than strictly necessary may generate an “over-decomposition” of conceptually unique sources of signal change into multiple sources, resulting in a poor representation of the phenomena. Fortunately, when the spatiotemporal structure of interesting fMRI sources in the data is compatible with the constraint of spatial independence, the number of principal components selected is not critical, and accurate and robust separation of these sources can be achieved using a reasonable amount of computational resources of time and memory.

For the extension of ICA from single to multi-subjects data sets, the three methods of applying the ICA model to group studies have been recently reviewed and compared on simulated data sets in (Schmithorst and Holland, 2004).

Group ICA by self-organizing clustering (sogICA)

After pre-processing of the fMRI time-series, ICA decomposition of each single-subject data set is performed using the FastICA algorithm (Hyvärinen, 1999). Alternatively, the Infomax algorithm can be used for the same purpose (Esposito et al., 2002). All data analysis steps after pre-processing and before brain map generation were performed in Matlab™ environment. Before entering the core ICA routines, a 3-D mask common to all subjects' data sets was defined in the Talairach space and applied to the image time-series in order to exclude the voxels outside the brain. Appropriate variables were defined in a Matlab™ workspace and the FastICA Toolbox 2.1 (<http://www.cis.hut.fi/research/software.shtml>) was used to set all the parameters of the decomposition, including dimensionality reduction, orthogonalization approach (symmetric or deflationary) and the nonlinearity. Using PCA, we reduced data sets from all the subjects and all the experiments to 40 temporal dimensions, and 40 ICs per subjects were estimated using the symmetric orthogonalization approach of the FastICA algorithm.

The ICA estimates from each subject were organized in one single set of components (maps in \mathbf{C} and time-courses in \mathbf{A} following the notations of Eq. (2)), with an additional label preserving the link from the components to the original subject; those components were then clustered according to their mutual similarities as follows.

A natural measure of similarity between the estimated independent components is the absolute value of their mutual correlation coefficients, in space for the source estimates or in time for the associated basis time-courses; a flexible similarity measure between component i and j $\mathbf{SM}(i, j)$ based on the classical Pearson correlation coefficient can be defined by combining the spatial and temporal correlations of the components in one formula:

$$\mathbf{SM}(i, j) = \lambda \cdot \mathbf{CCs}(i, j) + (1 - \lambda) \cdot \mathbf{C} \mathbf{Ct}(i, j) \quad (3)$$

$CCs(j,j)$ being the spatial correlation coefficient of independent components C_i and C_j and $CCT(i,j)$ being the temporal correlation between the two associated time-courses of activation of C_i and C_j , that is, the column vectors A_i and A_j of the corresponding mixing matrix estimate obtained from the ICA decompositions. The λ parameter is bounded between 0 and 1 and allows a user-defined interactive weighting of temporal or spatial similarity of the components: for instance $\lambda = 0.5$ implies an equal weighting of 50% spatial and 50% temporal similarity measure. In this study, we generated spatial ICA estimates and, thus, $\lambda = 1$ (pure correlation in the space domain) was used as default, but different (adaptive) choices for the similarity maybe justified in specific contexts (Kherif et al., 2003).

The $M \times M$ similarity matrix SM is, then, transformed into a dissimilarity matrix DM as in (Himberg et al., 2004):

$$DM(i,j) = \sqrt{1 - SM(i,j)} \quad \forall i,j = 1, \dots, M \quad (4)$$

This matrix is used as a “distance” matrix in the original space of the components and represents the input for the clustering step.

In this work, we used an approach to clustering that exploits the natural organization of components in the subject space (self-organizing group ICA, sogICA) by using the inter-subject variability in space of the components as the main dimension of investigation of the similarity matrix.

A supervised hierarchical clustering algorithm, linking the components to each other only when differently labeled (i.e., belonging to different subjects), was implemented as follows:

A threshold for the maximum “within-cluster” distance among components is initially set to zero; this threshold is then increased following a step-wise change from 0 to 1 (the step being as fine as the minimum distance between two different components in the whole space of components); at each step, a new cluster is detected when: (i) the mutual within-cluster distances are below the current value of the threshold and (ii) the cluster is “representative” of a group or subgroup of subjects according to the user’s specification about the minimum size of the group (with no repetitions). After cluster formation, the clustered components are excluded from successive steps. If the cluster includes a number of subjects lower than the total sample, similar components from subjects not yet included in the cluster may be integrated into it at later steps of clustering. Differently from typical clustering methods, there is not a check on the inter-cluster distance that is left unconstrained.

Following this approach, the user does not specify the number of clusters for the partition, but only the minimum number of subjects to be considered for a “group-representative” cluster: different searches can be, thus, performed by the user ranging from a default choice of including all the subjects of the study in a cluster, to a choice that aims to discover subgroups in the whole subject sample, either simply related or not to a given external factor.

As for ICA in general, an additional task that requires some form of interaction by the user is represented by the rank-ordering of the extracted clusters. The sogICA approach combines three natural ways of ranking the extracted clusters, through the minimum, mean and maximum within-cluster distances.

In order to help the user towards the interactive exploration of the “current level” of clustering, as well as to address interesting intra- and inter-cluster relations, a visual tool has been implemented and evaluated. Each IC estimate from each subject is plotted as a point on a graph, labeled according to the need of the group analysis (i.e., simply S1, S2, ... or M1, F1, M2, F2, ... for

males and females, etc.) and colored according to the cluster membership (i.e., one new color per new cluster).

The similarity matrix can be fruitfully visualized in a two-dimensional space using methods related to multidimensional scaling (MDS) (Torgerson, 1952). A typical MDS plot is obtained from the distance matrix in Eq. (4) and follows the fundamental criterion that the Euclidean distances of the points in the plot are to be as similar as possible to the original distances in the matrix DM . In other words, MDS methods look for an optimal configuration of points in a (for example) 2-D space by minimizing the mismatch between the distances between the points in the plot and the original distances in the matrix. MDS was previously used in brain imaging by (Friston et al., 1996) for the study of intra-individual functional connectivity and very recently proposed for the study of entire subjects’ data sets or model-driven activity patterns’ homogeneity in Kherif et al. (2003).

As in a previous application of MDS to neuroimaging time-series (Himberg et al., 2004), we used Curvilinear Component Analysis (CCA) (Demartines and Hérault, 1997) for cluster visualization because we have found this variant to be computationally efficient as well as suitable for this present application.

Traditional MDS methods (Mead, 1992) aim at faithfully preserving distances in the output by minimizing a cost function (“stress”) between the original and scaled distances. The cost function for CCA can be presented as:

$$E = \sum_i \sum_{j \neq i} \left[(\hat{d}(i,j) - d(i,j))^2 \cdot g(\hat{d}(i,j)) \right] \quad (5)$$

where $g(\cdot)$ is a function of projected distances $\hat{d}(i,j)$ and $d(i,j)$ are the distances in the original space. The cost function E in Eq. (5) with $g = 1$ defines a classical MDS method with the “raw stress” that aims at global preservation of distances. In the CCA, g is a bounded, monotonically decreasing function, for example, an inverse exponential function. Thus, the distances that are short in the projection have more weight in the stress of the CCA. This feature is expected to favor the local topology of points in the projection and makes CCA strictly related to self-organizing map (SOM) theory (Kohonen, 1989). Nonetheless, the CCA differs from SOM as it does not quantize the data and it is not confined to a regular grid in the output space.

Modified routines from the SOM toolbox 2.0 (Vesanto et al., 2000) for MatlabTM were used for the implementation of the MDS visualization (<http://www.cis.hut.fi/research/software.shtml>).

Since the original distances in DM are obtained from a huge number of dimensions, only a certain amount of the total variance contained in the high-dimensional similarity matrix will be correctly represented in the two-dimensional configuration, and geometrical distortions will be likely to occur in the plot. Nonetheless, when the points are plotted in a Euclidean subspace, the resulting distribution will tend to capture, although in a parsimonious way, the main structure of the comparative similarities (Friston et al., 1996). As a consequence, it is to be expected that the generated plot will be a trustworthy representation of the components’ similarities at least at the highest level of the clustering exploration when the corresponding entries of the similarity matrix explain most of the variance of the similarity measure. As far as the resulting distribution of points will capture this structure, the MDS plot display will function as an interactive visual tool for the evaluation and selection of group-representative clusters: relevant clusters in the original high-dimensional and

more informative space of the components are, thus, expected to be projected as geometrically consistent clusters of points in the low (2)-dimensional space, while less dense cluster geometries may suggest further investigation of the actual homogeneity of the components' structure across the subjects included in the cluster.

The clustering procedure as well as the MDS projection can eventually be iterated many times by changing the minimum number of subjects required to define the cluster, in order to discover possible subgroups.

Once the user has retrieved the estimates belonging to a cluster, the average component of this cluster is computed and, henceforth, assumed as the group component representative of the cluster. A random-effects group component can be easily generated by dividing voxel-wise the mean and the variance across subjects of the clusters' component values (Calhoun et al., 2001b).

Definitely, the proposed approach, as an ICA method that produces a proper sum of the separated independent components for the imaging result of fMRI data processing, can be put in the context of hybrid or alternative ICA approaches (Chen and Yao, 2004; McKeown, 2000).

Results

Fig. 2 shows the plots of the minimum, mean and maximum (spatial) similarity distances for the six subjects clusters obtained from the 240 components extracted in each of the two experiments. All six subjects contributed to all the clusters. For these plots, the distances have been ordered according the minimal distance, but the mean and the maximum distances are shown as well to provide a general description of the quality of all the clusters. In the following, we show the graphs, the maps and the time-courses corresponding to the first six clusters according to the minimum intra-cluster distance, in order to highlight the tightest structures in the performed clustering.

Fig. 3 shows the cluster plot for experiment 1 as it appears for the first six clusters selected. For all six clusters, except cluster 1, the random-effects group components have been projected onto the inflated cortical surfaces of the MNI template, thresholded at $z = 1.65$ ($P < 0.05$) with a color code for the activated vertices corresponding to the color of the cluster in the plot. Fig. 4 shows the same five clusters (from 2 to 6) in terms of individual inflated

maps. A relatively low threshold was used to produce a more comprehensive view of the spatial distribution of the components. In order to facilitate the 3-D inspection of the areas of activation, bottom (occipital) views are presented for the "visual" components (cluster 2 and 4), parietal views for clusters 2, 3, and 5 and a frontal view for cluster 6. The average time-course of each cluster is plotted and reported near to the maps in Fig. 3 and (in bold) superimposed on the individual component time-courses in Fig. 5.

Cluster 2 is easily classified as the expected consistently task-related (CTR) component: as expected a large bilateral region is active in this component, including the primary retinotopic areas in all six subjects (Sereno et al., 1995); the time-course of activity is highly correlated to the task. The other four clusters exhibit a high (3 and 5) to intermediate (4) or negligible (6) correlation with the task, and can be associated, respectively, to different spatially independent processes that are transiently or non-task related.

The random effects map of cluster 4 shows a bilateral region of activation in the medial part of the occipital lobes. The time-course is transiently task-related (TTR), with most of the signal change time-locked to the onset and the off-set (transitions) of the visual stimulation. The spatial layout of the activation map resembles a foveal representation of the visual stimulus as presented in (Tootell et al., 1997). The presence of this source may be related to the stable presence of the fixation cross that causes the perception of abrupt eccentricity change in the retinal stimulation. Similar effects were also reported in previous ICA studies with similar fMRI visual paradigms (Calhoun et al., 2001b, 2003).

Clusters 3 and 5 exhibit minimally overlapping posterior cingulate cortical (PCC) activity of different extensions: in cluster 3, a smaller PCC activity is accompanied by an inferior parietal bilateral activity, possibly related to parietal eye field (PEF) activity (even if located inferior than usually reported (Pierrot-Deseilligny et al., 2004)), while in cluster 5, the layout of the component is more posterior and wider. Overall, the layout of cluster 3 is highly reminiscent of a spatial pattern previously found in passive (resting) state experiments (Raichle et al., 2001). The different contribution of this inferior parietal activity to clusters 3 and 5 is also consistent and evident in the individual maps of Fig. 4. Lateral spots are also present in cluster 5 (see the individual maps in Fig. 4), but location and laterality are less spatially stable and consistent across subjects (this is also somehow suggested by the more extended geometry of cluster 5 in the MDS plot). Since the eye




Fig. 2. Plots of the intra-cluster similarity measure (minimum, average and maximum) as a function of the cluster index after re-ordering on the basis of the minimum intra-cluster measure. All six subjects contribute to all the clusters. Left: Experiment 1. Right: Experiment 2.

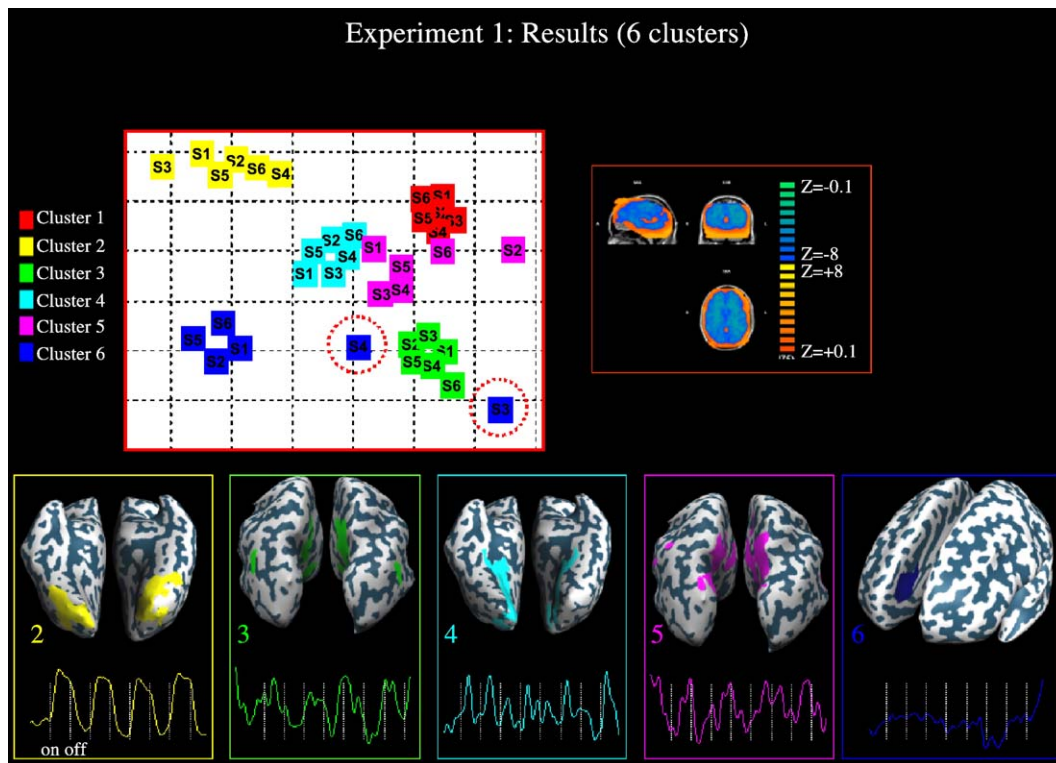


Fig. 3. Cluster plot as obtained by MDS projection using CCA for the ICA decomposition in Experiment 1 with the first 6 clusters displayed. The random-effects activation maps ($z = 1.65$) and the average (normalized) time-courses are reported with colors according to the cluster plot. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

movement activity was not externally controlled, the interpretation of these clusters can only be hypothetical. The detected functional connectivity patterns may nonetheless be ascribed to task-induced eye movements or attention and impulse control distributed processes related to the visual task and known to be served by multiple cingulate regions (Peterson et al., 1999). In cluster 5, the limited spatial consistency of the parietal components hides them in the group representation.

The average time-course of cluster 6 is poorly task-related and dominated by low temporal frequency components, and suggests the interpretation of the source process as a much slower modulation effect of the baseline activity: the main focus of the activity of cluster 6 is located again in a midline cingulate region, specifically in the ventral anterior cingulate cortex (vACC), and its signal is slightly negatively related to the task. Despite the low correlation with task and the high intersubject variability (Duann et al., 2002) of the component time-courses (see Fig. 5), the individual spatial patterns exhibited a surprisingly high regularity across all the subjects of the study (Fig. 4). Due to its spatial distribution, this vACC activity is likely to be ascribed to attention shifts, rather than to the so-called default mode of attention-related brain function (Raichle et al., 2001). Nevertheless, the vACC activity exhibits a nice degree of spatial similarity with the vACC activity reported by Greicius et al. (2003) in a resting state study. For cluster 6, the MDS plot in Fig. 3 shows a dislocation of the ICs from subjects S3 and S4, and, thus, a possible deviating structure of the patterns: a detailed inspection of the individual maps confirms a slight dislocation of the ACC activity foci for S3 and S4 in terms of the z coordinate in the Talairach space (more superior foci in S3 and S4 and more inferior in S1, S2, S5, and S6).

Despite the highest spatial intra-cluster similarity, all the components belonging to cluster 1 exhibited the least “focal” spatial structure in their respective individual decomposition. A tri-planar slice projection of positive and negative values of this cluster is presented in Fig. 3: the color-scaled representation between $z = 0$ and $z = 8$ highlights the special structure of its spatial distribution over the whole brain mask, with voxels of the same tissue type (for instance white matter) preferentially grouped together at the highest values of the component. A similar cluster of components was consistently obtained in experiment 2 (see Figs. 6 and 7) as well: for each individual, the estimated kurtosis of these component always corresponded to the lowest value among all the component kurtosis values although they always contributed for the maximum percentage of the total variance explained by all the estimated components. In order to rule out the possibility that cluster 1 was generated by a failure of the FastICA algorithm to converge sufficiently, we repeated the ICA estimation step on the same data sets by using the infomax algorithm, but found again the “cluster 1” with the same general character. These observations suggest that these components are not due to a problem of estimation but rather may represent a sort of “residual” component implied by the same ICA model. In fact, since the basic ICA model attempts to explain and decompose all the variance, the sources that, because of their distribution, are less fitted to the statistical model will tend to stay confined in one single (residual) component, with a low-kurtosis map. Interestingly, components with a structure similar to that of cluster 1, roughly discriminating voxels on the basis of an average tissue-specific signal change, were extracted from perfusion MRI time-series by Kao et al. (2003). The high degree of similarity across subjects of these




Fig. 4. Individual maps for the five of the six clusters (2–6) selected from Experiment 1.

components that make them robustly extracted in sogICA may be ascribed to the smoothing, normalization, and masking procedures that are common to all of the subject data sets that undergo the FastICA decomposition, and that “shape” both signal and noise of the data set in the same anatomical space.

In Fig. 6, we report the impact on the subject-level “grouping” of the components of the injection of additive Gaussian noise (AGN) in one subject’s data set (S1) from experiment 1. Specifically, we report the mean intra-cluster similarity distance, for the CTR and (after averaging) for a different number of clusters, at different levels (standard deviation) of injected AGN (25%, 50%, and 100%). In this simulation, as in McKeown et al. (1998a,b), we assumed the CTR component as signal (“truth”) and estimated a “baseline” noise from the time-courses of the voxels that fall within the lowest quartile of the distribution. As also McKeown et al. (1998a,b) did, we observed that the individual decompositions of the modified data sets were highly robust with respect to AGN and show in Fig. 6 how the slight loss of accuracy observed in the CTR component of subject S1 did not significantly affect the subject-level clustering of all the CTR components.

Fig. 7 shows the MDS plot and 6 clusters from experiment 2. The cortical surface projections of the maps, thresholded at $z = 1.65$ are shown with the color code of the activity corresponding to the colors of the clusters in the MDS plot.

Clusters 5 and 6 correspond to the expected CTR components for this paradigm (left and right hemifield stimulation, respectively). The inspection of the midline brain activity allows to

recognize that clusters 3 and 4 in experiment 2 correspond, respectively, to clusters 5 and 3 in experiment 1 in Fig. 3, even if lateral spots in cluster 4 are less symmetric in their bilateral locations. Cluster 2 in experiment 2 corresponds to cluster 4 in experiment 1.

In order to test the sensitivity of the method to an increased intersubject variability, we considered the interesting case when one subject does not have one component that the other subjects of the groups have (i.e., this subject did not perform one of the two tasks for some reason). To this purpose, we repeated the decomposition of subject S1 after having artificially modified the input data set in order to alter as much as possible one single component (namely the CTR components belonging to cluster 6 in Fig. 7). This was done by injecting Gaussian noise in the data of standard deviation equal to the that of the original signals in those voxels that were maximally active in the chosen component ($z > 5$): this way, we substantially affected the spatial layout of the chosen component without altering the layout of the other extracted components whose patterns remained practically the same compared with the previous decomposition on original data. Fig. 8 reports the result of the new clustering and the plot for a new CCA-based MDS projection of the estimates. All the clusters are practically identical to Fig. 7, with a slight difference in the layout of cluster 6 where the spatial power of the component from subject S1 was artificially reduced. The result of the projection clearly warns against the possible inhomogeneity of the group caused by the deviating component of subject S1. A further deviation of




Fig. 5. Individual component time-courses for the six clusters selected from Experiment 1. The mean time-courses, as representative time-courses of the cluster components, are shown in bold superimposed on the individual time-courses.

subject S3 in cluster 2 is noted as well that was also present in the projection of Fig. 7.

Discussion

We have presented a new method for the extraction and evaluation of fMRI group activation maps through the application of ICA to single-subject fMRI data. We have applied it to an illustrative visual paradigm in order to verify the capability of the method to extract blindly meaningful patterns of brain activation in an fMRI group study. For this purpose, the natural, self-organizing tendency of individual components to form clusters in the subject space according to simple measures of mutual similarity of their spatiotemporal structures was utilized. In this work, the spatial correlation was used for the analysis, although a general combined measure is also suggested for those experimental settings where more temporal synchrony across subjects is to be expected (Bartels and Zeki, 2004).

In all the experiments, we were always able to blindly identify one (full field stimulation) or two (hemifield stimulations) consistently task-related components activating the primary visual areas in the occipital lobe in all six subjects, resulting in sufficiently accurate random-effects group maps.

A bilateral transiently task-related visual component was identified in both experiments and in all the subjects based solely on the high level of inter-subject similarity of the corresponding

pattern. A similar component was already reported in previous ICA studies with a similar paradigm (Calhoun et al., 2001c, 2003) and was related to the transitions of the visual paradigm by inspection of the time-course of activation.

The cluster plots obtained by MDS projections of the individual component estimates proved to be a useful tool to address differences in the individual components that the algorithm blindly classifies as homologous components according to the similarity measures. Although affected by the inevitable distortion of using a very low number of dimensions for the geometrical representation, the MDS projection provided by CCA exhibited a fairly stable performance and its display with the labeling according to the cluster membership proved to capture a reasonable amount of the variance (2.3% for experiment 1 and 2.7% for experiment 2) in a way that intercluster similarity as well as within cluster outliers or deviating components could be detected. The check for the homogeneity of the cluster is a useful step to improve the cluster representative maps by using subgroups and can represent a preliminary step for the ICA decomposition of aggregated data sets restricted to a more homogenous sample of data sets. The choice of presenting only six clusters was exclusively due to display purposes. Limiting the visualization to the “first” six clusters, corresponding to the six best ranked clusters in terms of “group representativeness”, allowed a quick inspection of the graph in conjunction with the associated component patterns. More clusters may be displayed and presented by the user, with the sole limitation that each newly appearing cluster, whose members are

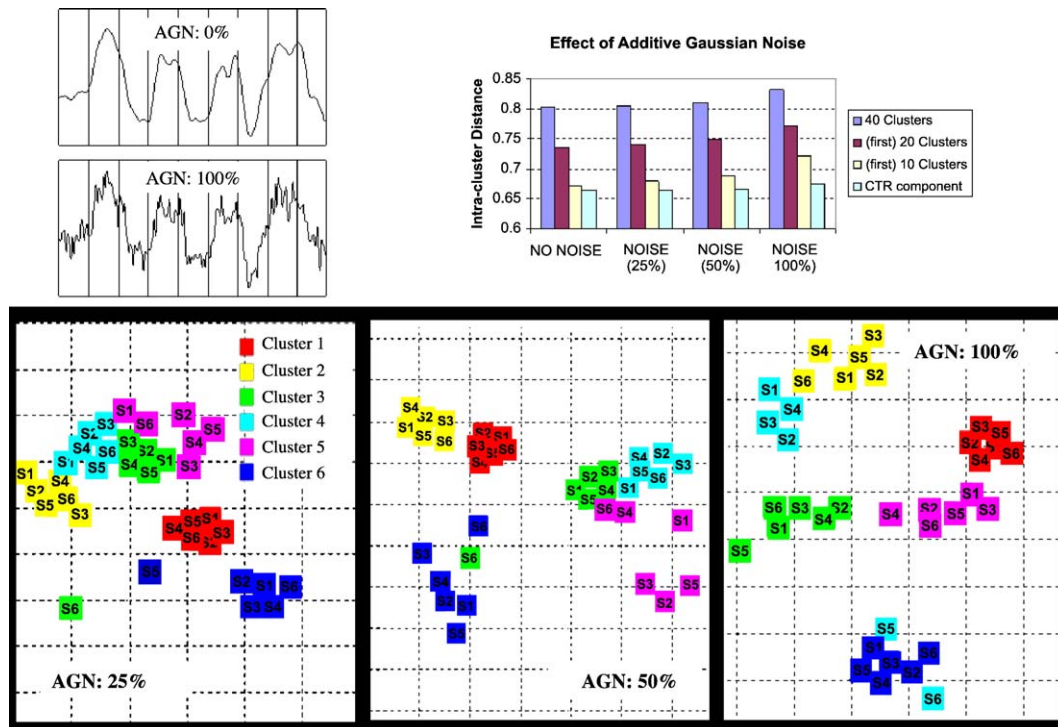


Fig. 6. Robustness of sogICA to additive Gaussian noise (AGN). Upper panel: (left) effect of AGN injection on a task-related signal is illustrated; (right) comparison in terms of mean intra-cluster similarity distances of the results at the investigated levels of AGN. Lower panel: cluster plots for three case of injected AGN in subject S1.

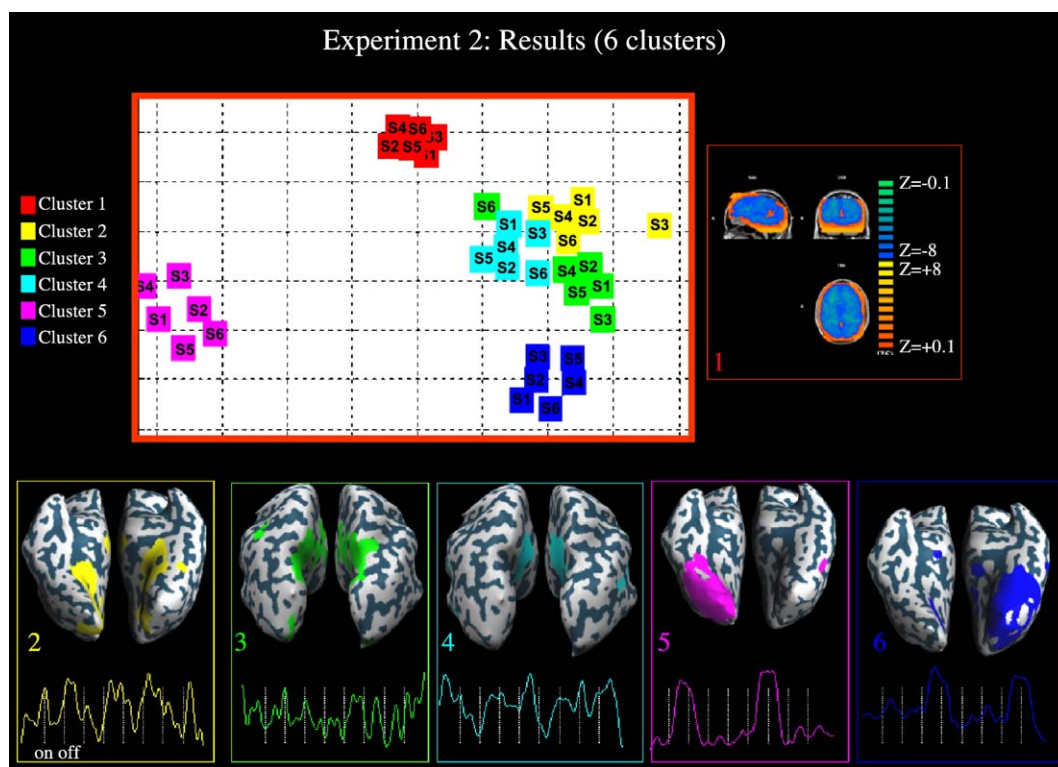


Fig. 7. Cluster plot as obtained by MDS projection using CCA for the ICA decomposition in Experiment 2 with the first 6 clusters displayed. The random-effects activation maps ($z = 1.65$) and the average (normalized) time-courses are reported with colors according the cluster plot. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

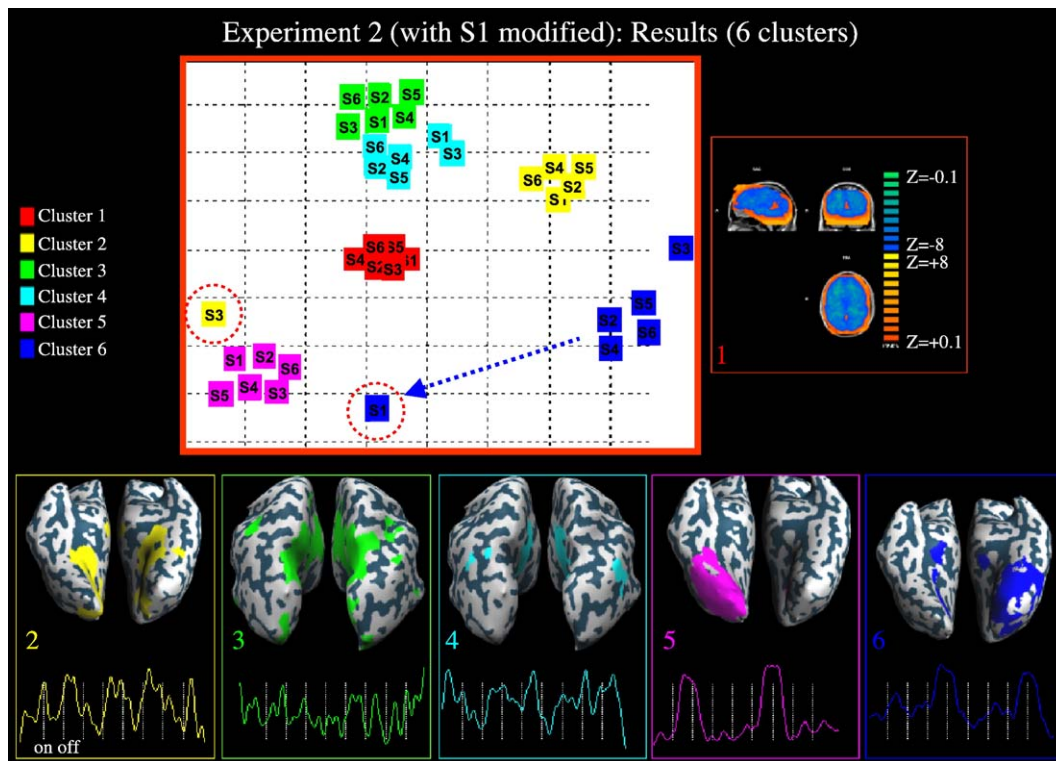


Fig. 8. Cluster plot as obtained by MDS projection using CCA for the ICA decomposition in Experiment 2 after the data set from S1 was modified by reducing the spatial power of one of the two CTR components (the blue one). The first 6 clusters are displayed. The random-effects activation maps ($z = 1.65$) and the average (normalized) time-courses are reported with colors according to the cluster plot. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

expected to be more and more widespread on the plane of projection, will make the visualization more difficult to interpret.

The visual inspection of both temporal waveforms and spatial maps suggests the success of the method in grouping the correct IC estimates from multiple subjects for the expected sources of bilateral visual activation. Thus, not only the method represents a valid stand-alone alternative to the ICA approaches based on aggregate group data sets, but it is also capable of selecting blindly those interesting patterns that form relevant clusters in the subject space. In this sense, the proposed approach also acts as a tool for the rank-ordering and the parsimonious selection of those components that are relevant for the group: using the intersubject similarity for the selection does not require any assumption about the expectability of temporal responses or the anatomical layouts of the components. Those aspects render the sogICA framework specifically interesting for complex naturalistic fMRI settings like those presented by Hasson et al. (2004) and Bartels and Zeki (2004). In these studies, outstanding evidence has been produced of the striking level of voxel-by-voxel activity synchronization existing between different human brains and of how the resulting subject-to-subject similarity can be the crucial factor for the generation of meaningful fMRI group patterns.

Other applications may be of interest, when there is partial knowledge or even a simple suspicion about possible clustering of the sampled subjects with respect to some external analogue measure (e.g., performance or other phenotypic measures) for which the classification in discrete groups is not trivial. For what concerns the use of genotypic information with brain imaging data (Hariri et al., 2002), reverse statistical analysis from subgroup

patterns to specific polymorphisms may be potentially addressed by the investigation of the self-organizing clustering of the components. In fact, the user may search for specific cluster geometries that suggest a pattern-based (blind) grouping and, thus, a subject classification, a nice property for an improved characterization of the differences between a population of patients and healthy subjects.

In principle, both the similarity measure and the clustering method can be flexibly and freely improved in future development of the sogICA framework, possibly according to the type of data or the experimental design: different types of group inference may be addressed by different measures of mutual similarity that exploit single or multiple combined features of components. In this introductory paper, we used the most natural and general measure of similarity between the estimated independent components, the absolute value of their mutual spatial correlation coefficients, and suggested the use an MDS-based plot to supervise the clustering of the components: the cluster analysis through the self-organizing mapping of the components in a simple bidimensional plot enables the use of the spatial proximity of the components as a rapid, although approximate, visual estimate of similarity. The clustering of the estimated components is expected to yield two kinds of crucial information: first, the presence of homologous sources of brain activity in different subjects that possibly code for a robust experimental factor; second, the degree of relative variability of this effect across the selected subjects and the possible confounding effects that interacts with the identified and selected ones. A compact, tight cluster emerges when a given source is robustly extracted in different decompositions from different subjects,

suggesting the presence of a similar phenomenon in all the subjects of the sample. More sparse clusters will suggest effects that are more variably measured in the sample; moreover, multiple adjacent disjointed clusters may suggest “subgroup” association or definition for a further between-group analysis step.

In conclusion, this study demonstrates the feasibility and addresses the potential of using single-subject ICA estimates for an fMRI group analysis by exploration of the self-organizing clustering of individual components in the subject space. Since we used the conventional ICA generative model for single-subject data sets and already validated algorithms for ICA estimation, we did not use simulated activation data but tested the approach using real activation experiments with an experimental design ensuring one or two concurrent sources of brain activity. More complex experimental settings as well as a tailored simulation framework will improve the method in future application and method studies.

Acknowledgments

Study was supported by Swiss National Science Foundation (grant no. PP00B-103012) and by Academy of Finland (project #106473).

References

- Bartels, A., Zeki, S., 2004. The chronoarchitecture of the human brain—natural viewing conditions reveal a time-based anatomy of the brain. *NeuroImage* 22, 419–433.
- Bell, A.J., Sejnowski, T.J., 1995. An Information-Maximisation approach to blind separation and blind deconvolution. *Neural Comput.* 7, 1004–1034.
- Brown, G.D., Yamada, S., Sejnowski, T.J., 2001. Independent component analysis at neural cocktail party. *Trends Neurosci.* 24, 54–63.
- Burbaud, P., Camus, O., Guehl, D., Bioulac, B., Caille, J., Allard, M., 2000. Influence of cognitive strategies on the pattern of cortical activation during mental subtraction. A functional imaging study in human subjects. *Neurosci. Lett.* 287, 76–80.
- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001a. Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Hum. Brain Mapp.* 13, 43–53.
- Calhoun, V.D., Adali, T., McGinty, V.B., Pekar, J.J., Watson, T.D., Pearlson, G.D., 2001b. fMRI activation in a visual-perception task: network of areas detected using the general linear model and independent components analysis. *NeuroImage* 14, 1080–1088.
- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001c. A method for making group inferences from functional MRI data using independent component analysis. *Hum. Brain Mapp.* 14, 140–151.
- Calhoun, V.D., Adali, T., Pearlson, G., Pekar, J., 2001d. Group ICA of functional MRI data: separability, stationarity and inference. *Proc. Int. Conf. ICA and BSS*, 155–160.
- Calhoun, V.D., Adali, T., Pekar, J.J., Pearlson, G.D., 2003. Latency (in)sensitive ICA. Group independent component analysis of fMRI data in the temporal frequency domain. *NeuroImage* 20, 1661–1669.
- Calhoun, V.D., Pekar, J.J., Pearlson, G.D., 2004. Alcohol intoxication effects on simulated driving: exploring alcohol-dose effects on brain activation using functional MRI. *Neuropsychopharmacology* 29 (11), 2097–2117.
- Castelo-Branco, M., Formisano, E., Backes, W., Zanella, F., Neuenschwander, S., Singer, W., Goebel, R., 2002. Activity patterns in human motion-sensitive areas depend on the interpretation of global motion. *Proc. Natl. Acad. Sci. U. S. A.* 99, 13914–13919.
- Chen, H., Yao, D., 2004. Discussion on the choice of separated components in fMRI data analysis by spatial independent component analysis. *Magn. Reson. Imaging* 22, 827–833.
- Demartines, P., Herault, J., 1997. Curvilinear component analysis: a self-organizing neural network for nonlinear mapping of data sets. *IEEE Trans. Neural Netw.* 8, 148–154.
- Duann, J.R., Jung, T.P., Kuo, W.J., Yeh, T.C., Makeig, S., Hsieh, J.C., Sejnowski, T.J., 2002. Single-trial variability in event-related BOLD signals. *NeuroImage* 15 (4), 823–835.
- Esposito, F., Formisano, E., Seifritz, E., Goebel, R., Morrone, R., Tedeschi, G., Di Salle, F., 2002. Spatial independent component analysis of functional MRI time-series: to what extent do results depend on the algorithm used? *Hum. Brain Mapp.* 16, 146–157.
- Esposito, F., Seifritz, E., Formisano, E., Morrone, R., Scarabino, T., Tedeschi, G., Cirillo, S., Goebel, R., Di Salle, F., 2003. Real-time independent component analysis of fMRI time-series. *NeuroImage* 20, 2209–2224.
- Friston, K.J., Frith, C., Liddle, P., Frackowiak, R.S.J., 1993. Functional connectivity: the principal component analysis of large data sets. *J. Cereb. Blood Flow Metab.* 13, 5–14.
- Friston, K.J., Frith, C.D., Fletcher, P., Liddle, P.F., Frackowiak, R.S., 1996. Functional topography: multidimensional scaling and functional connectivity in the brain. *Cereb. Cortex* 6, 156–164.
- Friston, K.J., Holmes, A.P., Price, C.J., Buchel, C., Worsley, K.J., 1999. Multisubject fMRI studies and conjunction analyses. *NeuroImage* 10, 385–396.
- Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc. Natl. Acad. Sci. U. S. A.* 100, 253–258.
- Hariri, A.R., Mattay, V.S., Tessitore, A., Kolachana, B., Fera, F., Goldman, D., Egan, M.F., Weinberger, D.R., 2002. Serotonin transporter genetic variation and the response of the human amygdala. *Science* 297, 400–403.
- Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., Malach, R., 2004. Intersubject synchronization of cortical activity during natural vision. *Science* 303, 1634–1640.
- Himberg, J., Hyvarinen, A., Esposito, F., 2004. Validating the independent components of neuroimaging time series via clustering and visualization. *NeuroImage* 22 (3), 1214–1222.
- Hyvärinen, A., 1999. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans. Neural Netw.* 10, 626–634.
- Hyvarinen, A., Karhunen, J., Oja, E., 2001. *Independent Component Analysis*. Wiley.
- Kao, Y.H., Guo, W.Y., Wu, Y.T., Liu, K.C., Chai, W.Y., Lin, C.Y., Hwang, Y.S., Jy-Kang Liou, A., Wu, H.M., Cheng, H.C., Yeh, T.C., Hsieh, J.C., Mu Huo Teng, M., 2003. Hemodynamic segmentation of MR brain perfusion images using independent component analysis, thresholding, and Bayesian estimation. *Magn. Reson. Med.* 49, 885–894.
- Kherif, F., Poline, J.B., Meriaux, S., Benali, H., Flandin, G., Brett, M., 2003. Group analysis in functional neuroimaging: selecting subjects using similarity measures. *NeuroImage* 20, 2197–2208.
- Kohonen, T., 1989. *Self-Organization and Associative Memory*, third ed. Springer-Verlag, Berlin.
- Liao, R., Krolik, J., McKeown, M.J., 2004. Isolation and minimization of effects of motion on fMRI using multiple reference images. *Proc. ISBI*, 916–919.
- McKeown, M.J., 2000. Detection of consistently task-related activations in fMRI data with hybrid independent component analysis. *NeuroImage* 11, 24–35.
- McKeown, M.J., Jung, T.P., Makeig, S., Brown, G., Kindermann, S.S., Lee, T.W., Sejnowski, T.J., 1998a. Spatially independent activity patterns in functional MRI data during the stroop color-naming task. *Proc. Natl. Acad. Sci. U. S. A.* 95, 803–810.
- McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.-P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J., 1998b. Analysis of fMRI data by blind

- separation into spatial independent component analysis. *Hum. Brain Mapp.* 6, 160–188.
- McKeown, M.J., Hansen, L.K., Sejnowski, T.J., 2003. Independent component analysis of functional MRI: what is signal and what is noise? *Curr. Opin. Neurobiol.* 13, 620–629.
- Mead, A., 1992. Review of the development of multidimensional scaling methods. *Statistician* 41, 27–39.
- Peterson, B.S., Skudlarski, P., Gatenby, J.C., Zhang, H., Anderson, A.W., Gore, J.C., 1999. An fMRI study of Stroop word-color interference: evidence for cingulate subregions subserving multiple distributed attentional systems. *Biol. Psychiatry* 45, 1237–1258.
- Pierrot-Deseilligny, C., Milea, D., Muri, R.M., 2004. Eye movement control by the cerebral cortex. *Curr. Opin. Neurol.* 17, 17–25.
- Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., Shulman, G.L., 2001. A default mode of brain function. *Proc. Natl. Acad. Sci. U. S. A.* 98, 676–682.
- Schmithorst, V., Holland, S.K., 2004. Comparison of three methods for generating group statistical inferences from independent component analysis of functional magnetic resonance imaging data. *J. Magn. Reson. Imaging* 19, 365–368.
- Seifritz, E., Esposito, F., Hennel, F., Mustovic, H., Neuhoﬀ, J.G., Bilecen, D., Tedeschi, G., Scheﬄer, K., Di Salle, F., 2002. Spatiotemporal pattern of neural processing in the human auditory cortex. *Science* 297, 1706–1708.
- Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., Tootell, R.B., 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268, 889–893.
- Svensen, M., Kruggel, F., Benali, H., 2002. ICA of fMRI group study data. *NeuroImage* 16, 551–563.
- Talairach, J., Tournoux, P., 1988. *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme, Stuttgart.
- Tootell, R.B., Mendola, J.D., Hadjikhani, N.K., Ledden, P.J., Liu, A.K., Reppas, J.B., Sereno, M.I., Dale, A.M., 1997. Functional analysis of V3A and related areas in human visual cortex. *J. Neurosci.* 17, 7060–7078.
- Torgerson, W., 1952. Multidimensional scaling I—Theory and methods. *Psychometrika* 17, 401–419.
- Vesanto, J., Himberg, J., Alhoniemi, E., Parhankangas, J., 2000. SOM toolbox for Matlab 5. Report A57. Helsinki University of Technology, Neural Networks Research Centre, Espoo, Finland.