

## A group model for stable multi-subject ICA on fMRI datasets

G. Varoquaux<sup>a,c,\*</sup>, S. Sadaghiani<sup>b,c</sup>, P. Pinel<sup>b,c</sup>, A. Kleinschmidt<sup>b,c</sup>, J.B. Poline<sup>c</sup>, B. Thirion<sup>a,c</sup>

<sup>a</sup> Parietal project team, INRIA, Saclay-Île-de-France, Saclay, France

<sup>b</sup> INSERM unité 562, NeuroSpin, Saclay, France

<sup>c</sup> CEA, DSV, I2BM, Neurospin, Saclay, France<sup>1</sup>

### ARTICLE INFO

#### Article history:

Received 9 November 2009

Revised 2 February 2010

Accepted 4 February 2010

Available online 12 February 2010

### ABSTRACT

Spatial Independent Component Analysis (ICA) is an increasingly used data-driven method to analyze functional Magnetic Resonance Imaging (fMRI) data. To date, it has been used to extract sets of mutually correlated brain regions without prior information on the time course of these regions. Some of these sets of regions, interpreted as functional networks, have recently been used to provide markers of brain diseases and open the road to paradigm-free population comparisons. Such group studies raise the question of modeling subject variability within ICA: how can the patterns representative of a group be modeled and estimated via ICA for reliable inter-group comparisons?

In this paper, we propose a hierarchical model for patterns in multi-subject fMRI datasets, akin to mixed-effect group models used in linear-model-based analysis. We introduce an estimation procedure, CanICA (Canonical ICA), based on i) probabilistic dimension reduction of the individual data, ii) canonical correlation analysis to identify a data subspace common to the group iii) ICA-based pattern extraction. In addition, we introduce a procedure based on cross-validation to quantify the stability of ICA patterns at the level of the group. We compare our method with state-of-the-art multi-subject fMRI ICA methods and show that the features extracted using our procedure are more reproducible at the group level on two datasets of 12 healthy controls: a resting-state and a functional localizer study.

© 2010 Elsevier Inc. All rights reserved.

### Introduction

Much of our understanding of brain function gained through functional Magnetic Resonance Imaging (fMRI) has been derived by correlating experimental signals of task-driven activation with external stimuli or events. Yet, the covariance structure of the fMRI signal holds as much information as the paradigm. Indeed at the scale of functional neuroimaging, positive correlations between distant regions induced by the coordination of distributed neuronal activity in a particular experimental context may provide a fundamental insight into brain function. Pioneered by Biswal et al. (1995), studies of paradigm-free functional connectivity investigate correlations in the BOLD (Blood Oxygen Level Dependent) time series. Such correlation studies have been decisive in identifying large functional networks of distant regions (Cordes et al., 2000; Greicius et al., 2003). It has been established that the functional connectivity signals observed in fMRI contains neuronal information beyond physiological or scanner noise (Laufs et al., 2003; Shmuel and Leopold, 2009), so that part of these

signals can be considered a trace of underlying neuronal activity. In addition, recent work suggests that correlations in the functional signal are shaped by the anatomical connectivity structure (Greicius et al., 2009; Honey et al., 2007; Skudlarski et al., 2008; van den Heuvel et al., 2009).

Analysis of the correlation structure of the BOLD signal reveals large-scale patterns of correlated activity (Biswal et al., 1995; Lowe et al., 1998) and is expected to help understand the subdivision of the brain into cognitive systems that have coherent activity across time. These systems may be labeled as *networks* assuming that they result from underlying brain connections (Bullmore and Sporns, 2009; Honey et al., 2007) and serve distinct functions (Fox and Raichle, 2007; Smith et al., 2009).

The study of brain function through functional-connectivity mapping can be carried out without requiring the subject to perform a specific task. The so-called *resting-state* protocols can be easily applied, and are especially useful to include impaired subjects in a multi-group analysis. The networks identified by such experiments can give insights into the mechanisms of brain diseases and their modifications in pathological situation can serve as biomarkers to aid in clinical diagnosis (Garrity et al., 2007; Greicius, 2008; Greicius et al., 2004; Mohammadi et al., 2009; Wang et al., 2006). Recently, Seeley et al. (2009) have shown that large-scale brain networks of co-activation are adequate neuro-physiological units to study the impact of neuro-degenerative diseases.

\* Corresponding author. Laboratoire de Neuro-Imagerie Assisté par Ordinateur, NeuroSpin CEA Saclay, Bt 145, 91191 Gif-sur-Yvette, France.

E-mail address: [gael.varoquaux@normalesup.org](mailto:gael.varoquaux@normalesup.org) (G. Varoquaux).

<sup>1</sup> Funding from INRIA-INSERM collaboration. The fMRI set was acquired in the context of the SPONTACT ANR project.

The study of the correlation structure of brain activity is plagued by the size of the data: an fMRI dataset comprises more than 10 000 time series associated with the various voxels, which yield millions of possible pair-wise correlations. Functional-connectivity analysis has been pioneered through seed-based studies (Biswal et al., 1995; Cordes et al., 2000; Fox et al., 2005) that potentially uncover large-scale networks of brain activity correlated with a user-specified seed region. This approach, although very successful, is limited by the prior choice of seed regions of interest. Various clustering techniques (Cordes et al., 2002; Golland et al., 2007; Thirion et al., 2006) have been used to automatically define the regions or signals of interest. However, spatial ICA (Kiviniemi et al., 2003; McKeown et al., 1998) is, to date, the most popular method for identifying meaningful patterns in correlation studies without prior definition of any target region. The patterns extracted by ICA are usually easy to interpret in a cognitive neuroscience context, as they are most often well-contrasted, indicate different underlying physiological, physical, and cognitive processes, and can often be related to networks observed in different contexts, such as in seed-based analysis, or cognitive networks known from the literature (Smith et al., 2009).

ICA extracts salient patterns that are embedded in the data, and are thus considered as important to describe it. It is a purely data-driven method based on a loosely constrained data model; as a consequence, statistical significance of the extracted patterns remains unclear. In particular, there is no simple way to extrapolate from findings obtained in one dataset to other datasets, even if these are sampled from the same population. In the context of group analysis, statistically well-controlled seed-based studies have shown that the BOLD signal contains patterns of correlation that are highly reproducible across subjects, including in resting-state experiments (Shehzad et al., 2009). Similarly, some patterns extracted from resting-state fMRI datasets by an exploratory ICA approach are consistent at the group level (Damoiseaux et al., 2006) and have been used as biomarkers for population comparison (Garrity et al., 2007; Sorg et al., 2007).

However, ICA patterns can be relatively sensitive to mild data variation. Various, often non-overlapping, ICA patterns of group-level coherent activity have been reported, from resting-state data (Beckmann et al., 2005; Damoiseaux et al., 2006; Kiviniemi et al., 2003; Perlberg et al., 2008), and in task-based experiments (Beckmann and Smith, 2005; Calhoun et al., 2001). Probabilistic models have been used to provide pattern-level noise-rejection criteria (Beckmann and Smith, 2004) or goodness-of-fit measures of the model (Guo and Pagnoni, 2008), but still they do not provide pattern-level significance testing. The uncontrolled variability of the individual patterns is detrimental to population studies: there is no established framework for between-group comparison or inference on ICA maps.

ICA being an exploratory analysis technique that estimates a mixing model specific to the data, it is not meaningful to compare directly patterns estimated on different individual subjects. On the contrary, group-level patterns can be specialized to each subject (Calhoun et al., 2001; Filippini et al., 2009). Different strategies have been adopted for group-level extraction of ICA patterns. Patterns estimated at the subject level can be *merged* to form group maps (Esposito et al., 2005; Perlberg et al., 2008) although this is a challenging task because the correspondence of individual maps may be hard to assess and the merging operation is difficult to model from a statistical point of view. Individual-subject volumes can be concatenated along the time axis to apply the ICA algorithm on the group data (Calhoun et al., 2001). Finally, Beckmann and Smith (2005) have developed a tensorial extension of ICA that estimates patterns across subjects sharing the same time course throughout the experiment.

In this paper, we present a novel group model for multivariate patterns in fMRI volumes and an associated estimation procedure to

extract group-level ICA maps modeling subject variability. The strength of this method called CanICA lies in the identification of a subspace of reproducible components across subjects using generalized canonical correlation analysis (CCA). Combined with an explicit noise model and resampling procedure, this enables automatic selection of the number of components. In addition, we introduce a cross-validation procedure and metrics to compare the stability of a set of multi-subject patterns across different sub-populations. We compare our method to state-of-the-art fMRI group ICA methods with different group models: concatenation and tensorial group ICA approaches. We do not compare to *merging* procedures since they do not rely on a linear model between individual subject-level datasets and group-level Independent Components (ICs) and thus cannot be formulated with a spatially resolved between-subject variability of group-level ICs. We show with cross-validation that features extracted by our method are more stable on a group of 12 controls, both in a resting-state experiment and in a traditional *activation detection* experiment with a known paradigm.

## Theory

### Spatial ICA model for fMRI data

ICA assumes that the observed data is the linear mixture of unknown base signals that are recovered based on measures of statistical independence. The underlying model is that of blind separation of independent sources:

$$\mathbf{B} = \mathbf{M}\mathbf{A}, \quad (1)$$

where the rows of the matrix  $\mathbf{B}$  are the observed patterns, and those of  $\mathbf{A}$  form the patterns corresponding to the estimated independent sources, and  $\mathbf{M}$  is a mixing matrix estimated by ICA. To set the notations, when considering spatial patterns or components, such as  $\mathbf{A}$  or  $\mathbf{B}$ , in this article, we will use  $n_{\text{patterns}} \times n_{\text{voxels}}$ -shaped matrices;  $n_{\text{patterns}}$  is a number that corresponds to the model order, or possibly the number of observations, depending on the context. Note however that the independence of the patterns extracted by ICA algorithms is not guaranteed in theory and rarely checked in practice, and that the independence criterion used to identify the components often boils down to sparse component extraction (Daubechies et al., 2009). There is no theoretical basis to consider that a pattern is representative of solely one independent process, for instance movement, although in practice ICA is so far one of the of factor analytic transformations (Langers, 2009) most suited to blind pattern extractions from fMRI data. In addition, it is not clear that, from a neuroscientific point of view, independence is the right concept to isolate brain networks, as no functional system is fully segregated.

The patterns  $\mathbf{B}$  present in the acquired fMRI volumes are confounded by observation noise. As a result, the ICA mixing model is most often applied on a subset of the acquired signal, after an initial data-reduction step. In most ICA methods, this step is carried out using a Principal Components Analysis (PCA), the order of which thus determines the dimension of the signal subspace and thus the number of sources extracted by ICA. In the context of fMRI data analysis, a probabilistic PCA model can be used to introduce a noise model as the basis for this subspace selection (Beckmann and Smith, 2004).

### Existing group models for ICA on multi-subject fMRI data

ICA is a multivariate analysis technique: voxel-based time courses are not characterized as such, but as part of signal fluctuations in the entire brain. As a consequence, the voxel-level group models used in standard mass-univariate analysis—random effects or mixed effects—cannot be applied directly to ICA patterns. Two main strategies have been used so far to extract group-level patterns for fMRI images.

A first approach, introduced by Calhoun et al. (2001), concatenates individual subject data and performs data reduction and ICA on the resulting dataset. Data reduction is done by PCA. Subject maps are obtained by applying the mixing model learned on the group to the data specific to a subject. The group model underlying the estimation procedure is that the images observed  $\mathbf{Y}$  in the individual datasets are generated by a mixture of the group-level ICA patterns  $\mathbf{A}$  with additional noise  $\mathbf{E}$ :

$$\mathbf{P}=\mathbf{M}\mathbf{A} \quad (2)$$

$$\mathbf{Y}=\mathbf{W}\mathbf{P}+\mathbf{E} \quad (3)$$

with  $\mathbf{Y}=[\mathbf{Y}_1^T, \dots, \mathbf{Y}_S^T]^T$  the observed individual subjects images concatenated along the time axis, and  $\mathbf{E}$  the noise.  $\mathbf{W}$  gives a subject-specific set of loadings and  $\mathbf{P}$  are the principal components spanning the inter-subject signal subspace. This model consists in the addition of a subject-dependent observation noise in Eq. (3) to the ICA model (Eq. (2) or Eq. (1)). The GIFT toolbox (Calhoun et al., 2001) (<http://icatb.sourceforge.net/>) and the MELODIC software used in *ConcatICA* mode (Smith et al., 2004) implement variations on this model. They differ in the way to represent the group-level patterns: GIFT uses a  $T$ -statistics-based thresholding on a random-effects analysis of the individually reconstructed maps  $\{\mathbf{M}^{-1}\mathbf{W}_i^{-1}\mathbf{Y}_i, i=1 \dots S\}$ . This amounts to building a posteriori a hierarchical model with two levels of variance. *ConcatICA*, on the other hand, directly thresholds the groups patterns (see further). In either case, the model is estimated in two steps, using principal component analysis (PCA) to solve Eq. (3), followed by a noiseless ICA algorithm to identify the independent components in Eq. (2). In the actual implementations of the estimation of the group model in Eq. (3), the generative model differs slightly from the model on an implementation basis, as data reduction steps may be implemented via successive PCAs to limit the numerical size of each step.

The tensorial extension to ICA developed by Beckmann and Smith (2005) uses a model inspired from PARAFAC (Harshman, 1970): the observed images  $\mathbf{Y}$  are modeled as a trilinear combination of group-level independent patterns  $\mathbf{A}$ , common time courses, and subject-specific loadings, with additional observation noise  $\mathbf{E}$ :

$$\mathbf{Y}=(\mathbf{M}|\otimes|\mathbf{N})\mathbf{A}+\mathbf{E}. \quad (4)$$

Matrices  $\mathbf{M}$  and  $\mathbf{N}$  give the mixing of independent patterns across subjects and components.  $(|\otimes|)$  is the Khatri-Rao product, which imposes a triple-outer-product relationship between subject components and group-level independent components: subject components share the same set of independent spatial patterns and time courses for each spatial map, but the contribution of each independent pattern varies from subject to subject. The hypothesis guiding the tensor ICA model is that, in a multi-subject experiment with external correlates, the time course corresponding to activation of cognitive networks is set by the experiment, and thus shared between subjects. On the other hand, tensor ICA may imply only a low degree of prevalence of the components with respect to the dataset at hand, because it allows arbitrary subject-level loadings on each of the independent components.

Guo and Pagnoni (2008) have described both approaches as special cases of a general decomposition model in which subject-specific images are a linear combination of group-level independent components:

$$\mathbf{Y}=\mathbf{M}\mathbf{A}+\mathbf{E} \quad (5)$$

with  $\mathbf{Y}=[\mathbf{Y}_1^T, \dots, \mathbf{Y}_S^T]^T$  the group data matrix made of the time-wise concatenation of observed individual subjects data. The Group-ICA model and the Tensor-ICA model can be seen as putting restrictions on the structure of the mixing matrix  $\mathbf{M}$ . Guo and Pagnoni (2008) propose a more general estimation algorithm of an unconstrained mixing matrix with an expectation maximization (EM) algorithm, to learn the group structure associated via the data. The limitation of this

EM approach is that it is based on modeling the histogram of the independent components with a mixture of an arbitrary number of Gaussian components.

#### Group comparison with ICA

While ICA is often used to compare functional connectivity between groups, the mixing model of ICA does not provide a natural statistical framework for comparing patterns estimated from different datasets, unlike the GLM framework. Two recent contributions have laid out statistical group-comparison procedures.

Guo and Pagnoni (2008) propose a goodness-of-fit measure for their mixture-of-Gaussian-based ICA generative model using an approximate likelihood ratio test to compare different mixing models and discriminate if a group is better modeled as a homogeneous population or as a set of different subgroups. However, this approach essentially assesses the amount of data variance fit by the model (which is measured by the likelihood ratio test) so that it may be systematically affected by dimension selection issues, while it assumes that the unmodeled variance is independent across voxels to allow tractable computations.

Rombouts et al. (2009) use Tensor-ICA and a two sample  $t$ -test on the loadings of the different group patterns between healthy controls and patients with dementia to identify the patterns represented unevenly in the two populations.

The above procedures provide important indications on data structure, but they do not directly highlight the difference between groups in the individual ICs. Some of these individual patterns represent biologically meaningful components and are used in cognitive studies and as biomarkers. There is thus a strong interest in basing group comparisons on the patterns themselves, rather than the complete mixing model, such as performed in Garrity et al. (2007), Greicius et al. (2004), and Mohammadi et al. (2009). The difficulty in comparing these patterns stems from the fact that the estimation performed by the ICA algorithm is not robust against mild data variation: some global differences may exist between two decompositions of datasets that resemble each other; for instance, an IC present in one dataset can show salient features that are separated into two components in another dataset, so that we may consider the IC as *split* into two in the other dataset.

## Materials and methods

### A multivariate extension of mixed-effects models


To achieve a better control on variability of the ICs due to individual difference between subjects, we introduce a generative model of the signal making the separation between subject-to-subject variability, and observation noise specific to a subject or a session. At the group level, we describe the BOLD signal by a set of patterns  $\mathbf{A}$  corresponding to different independent components. These different components are extracted in a signal space common to the group and spanned by principal components  $\mathbf{B}$ . The generative model relates these group-level patterns to the observed signal via different noise terms.

#### Generative model

*Group-variability model.* The activity recorded on each subject  $s$  can be described by a set of subject-specific spatial patterns  $\mathbf{P}_s$ , which are a combination of the group-level patterns  $\mathbf{B}$  and additional subject-variability:

$$\text{for each subject } s, \mathbf{P}_s=\mathbf{A}_s\mathbf{B}+\mathbf{R}_s, \quad (6)$$

with  $\mathbf{A}_s$  a loading matrix giving how much each pattern is represented in subject  $s$ , and  $\mathbf{R}_s$  a residual matrix giving the deviation from the



**Fig. 1.** Diagram of the estimation procedure: successive estimation steps are applied from left to right, starting from individual datasets, to end with group-level independent components.

group patterns. In other words, a group description can be written considering the group of patterns (vertically concatenated matrices)  $\mathbf{P} = \{\mathbf{P}_s\}$ ,  $\mathbf{R} = \{\mathbf{R}_s\}$ , and  $\mathbf{\Lambda} = \{\mathbf{\Lambda}_s\}$ ,  $s = 1 \dots S$ ,

$$\mathbf{P} = \mathbf{\Lambda} \mathbf{B} + \mathbf{R}. \quad (7)$$

$\mathbf{\Lambda}$  is a loading matrix relating the subject-level components for each subject to the group-level components. If there are  $n_{grp}$  components generating the signal at the group level, and each of the  $S$  individual datasets are described by  $n_{sbj}$ , the shape of the loading matrix  $\mathbf{\Lambda}$  is  $(S \cdot n_{sbj}, n_{grp})$  (see Fig. 1 for a diagram).

**Observation model.** For each acquisition-frame time point the observed data is a combination of different subject-specific patterns  $\mathbf{P}_s$  confounded by observation noise: let  $\mathbf{Y}_s$  be the resulting spatial images in BOLD MRI sequences for subject  $s$  (an  $n_{frames} \times n_{voxels}$  matrix),  $\mathbf{E}_s$  the observation noise, and  $\mathbf{W}_s$  a loading matrix such that:

$$\mathbf{Y}_s = \mathbf{W}_s \mathbf{P}_s + \mathbf{E}_s. \quad (8)$$

#### Parallel to mixed-effect models

The generative model can be written in a non-hierarchical form as:

$$\begin{bmatrix} \mathbf{Y}_1 \\ \vdots \\ \mathbf{Y}_s \end{bmatrix} = \begin{bmatrix} \mathbf{W}_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \mathbf{W}_s \end{bmatrix} \left( \begin{bmatrix} \mathbf{\Lambda}_1 \\ \vdots \\ \mathbf{\Lambda}_s \end{bmatrix} \mathbf{M} \mathbf{A} + \begin{bmatrix} \mathbf{R}_1 \\ \vdots \\ \mathbf{R}_s \end{bmatrix} \right) + \begin{bmatrix} \mathbf{E}_1 \\ \vdots \\ \mathbf{E}_s \end{bmatrix} \quad (9)$$

In the group ICA framework of Guo and Pagnoni (2008),

$$\mathbf{Y} = \tilde{\mathbf{M}} \mathbf{A} + \tilde{\mathbf{E}} \quad (10)$$

$$\text{with } \tilde{\mathbf{M}} = \begin{bmatrix} \mathbf{W}_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \mathbf{W}_s \end{bmatrix} \begin{bmatrix} \mathbf{\Lambda}_1 \\ \vdots \\ \mathbf{\Lambda}_s \end{bmatrix} \mathbf{M} \quad (11)$$

$$\text{and } \tilde{\mathbf{E}} = \begin{bmatrix} \mathbf{W}_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \mathbf{W}_s \end{bmatrix} \begin{bmatrix} \mathbf{R}_1 \\ \vdots \\ \mathbf{R}_s \end{bmatrix} + \begin{bmatrix} \mathbf{E}_1 \\ \vdots \\ \mathbf{E}_s \end{bmatrix} \quad (12)$$

In other words, the mixing matrix can be factored out in a subject-specific matrix and group-level matrix, and the rejected noise is the addition of two terms, a subject-specific one and a group-level one, that explain different components of the signal. If we compare the ICA analysis with a GLM analysis, as suggested in McKeown et al. (1998), in Eq. (10),  $\tilde{\mathbf{M}}$  corresponds to the design matrix, and  $\tilde{\mathbf{E}}$  to the residuals. We can see that the decomposition of these matrices in Eqs. (11) and (12) can also be identified as the non-hierarchical expression of corresponding terms in a mixed-effect model (see for instance Friston et al., 2005).

Thus, the two-level group model exposed above can be seen as a multivariate formulation of mixed-effects models, often used in a standard univariate GLM-based analysis, as it models two sources of variance. However, it is not hypothesis-driven and does not rely on the knowledge of external correlates.

#### Estimation procedure

In this section, we describe the estimation procedure used to extract the patterns of interest from the data. The different variables estimated corresponding to the generative model are noted with a hat: if  $\mathbf{A}$  is the population set of patterns introduced in the model,  $\hat{\mathbf{A}}$  denotes the corresponding patterns estimated from the data. Fig. 1 gives a diagram summarizing the model and the successive estimation steps.

#### Noise rejection using the generative model

We start from fMRI image sequences, for each subject, corrected for delay in slice acquisition and for motion, then registered to a common template space. We extract the time series corresponding to a mask of the brain, center and variance-normalize them, resulting in pattern matrices  $\{\mathbf{Y}_s, s = 1 \dots S\}$ , with  $s$  corresponding to the subject index. We separate reproducible group-level patterns from noise by estimating successively each step of the above hierarchical model.

**Subject-level identification of observation noise.** First, we separate observation noise  $\mathbf{E}_s$  from subject-specific patterns  $\mathbf{P}_s$  as in Eq. (8), through principal component analysis (PCA). The principal components explaining most of the variance for a given subject's data set form the patterns of interest, while the tail of the spectrum is considered as observation noise.

Specifically, for each subject, we use a singular value decomposition (SVD) of the individual data matrix,

$$\mathbf{Y}_s = \mathbf{U}_s \mathbf{\Sigma}_s \mathbf{V}_s. \quad (13)$$

We retain the first  $n_{sbj}(s)$  rows of  $\mathbf{V}_s$  to constitute the patterns  $\hat{\mathbf{P}}_s$  and the first  $n_{sbj}(s)$  columns of  $(\mathbf{U}_s \mathbf{\Sigma}_s)$  give their loadings  $\hat{\mathbf{W}}_s$  in the subject's data, that can be interpreted as time courses of the components. The residual constitutes the observation noise,  $\hat{\mathbf{E}}_s$ :

$$\hat{\mathbf{P}}_s = (\mathbf{V}_s)_{1 \dots n_{sbj}}, \quad (14)$$

$$\hat{\mathbf{W}}_s = (\mathbf{U}_s \mathbf{\Sigma}_s)_{1 \dots n_{sbj}}, \quad (15)$$

$$\hat{\mathbf{E}}_s = \mathbf{Y}_s - \hat{\mathbf{W}}_s \hat{\mathbf{P}}_s, \quad (16)$$

where  $n_{sbj}(s)$  is the number of extracted components describing the signal at the subject level for the subject  $s$ .

Setting the number  $n_{\text{sbj}}(s)$  of principal components to retain at the subject level is a difficult question and not the main focus of this article. In general, this number can be selected by analyzing the individual datasets. We have found that with information-theoretic criteria initially used for model-order selection in PCA algorithms (Beckmann and Smith, 2004; Calhoun et al., 2001; Minka, 2001), the number of identified sources is overestimated for long fMRI time series and many non-meaningful ICA components are extracted. Various studies have been conducted of the influence of PCA-model-order choice in fMRI ICA methods (Himberg et al., 2004; Kiviniemi et al., 2009; Li et al., 2007). The optimal number of ICs has been empirically found to lie between 20 and 60 components depending on the dataset.

We use the resampling method described by Mei et al. (2008) in the context of statistical shape modeling. This method estimates the number of principal components necessary to model the data in the presence of Gaussian-distributed noise by assessing the stability of the subspaces spanned by the first  $n_{\text{sbj}}(s)$  principal components. Mei et al. (2008) showed empirically that it does not lead to diverging model order when new datasets are generated from the selected principal components with additional Gaussian-distributed noise. For more details on the model-order selection algorithm and its robustness to noise, we refer the reader to Mei et al. (2008). We apply this algorithm on the individual subject datasets,  $\mathbf{Y}_s$ , after the various preprocessing steps, as it is input in our group model. For computational reasons, we estimate the subject-level number of components on only one subject's data if all the individual datasets share the same acquisition parameters and length and use this parameter for all individual datasets: we assume that  $n_{\text{sbj}}(s)$  is independent of  $s$ , and we write  $n_{\text{sbj}}$ .

*Canonical correlation analysis to estimate the group-level patterns.* At the group level, we are interested in identifying the sub-space common to each subject's patterns. For this purpose, we use generalized Canonical Correlation Analysis (CCA).

CCA is most often used to compare two multivariate datasets by finding the successive univariate components pairs of each dataset that maximize cross-correlation. The components pairs are called *canonical variables*, and the associated correlation is the *canonical correlation*. We use a generalization of CCA to multiple datasets (Kettenring, 1971; Krzanowski, 1979)—in our case, one dataset per subject. We start from the different whitened datasets<sup>2</sup>  $\hat{\mathbf{P}}_s$ , and concatenate them to form  $\hat{\mathbf{P}}$ . We perform an SVD on  $\hat{\mathbf{P}}$ :

$$\hat{\mathbf{P}} = \Upsilon \mathbf{Z} \Theta. \quad (17)$$

where  $\Upsilon$  and  $\Theta$  are rotation matrices, and  $\mathbf{Z}$  is the diagonal matrix of singular values. The rows of  $\Theta$  form the canonical variables, in other words the inter-subject reproducible components, and the singular values on the diagonal of the matrix  $\mathbf{Z}$  form the canonical correlations, which yield a measure of between-subject reproducibility. We retain the first  $n_{\text{grp}}$  vectors of  $\Upsilon$ , the canonical weights forming the patterns of interest at the group level:

$$\hat{\mathbf{B}} = (\Upsilon_s)_{1 \dots n_{\text{grp}}}, \quad (18)$$

We select the dimension  $n_{\text{grp}}$  by keeping only the canonical variables for which the corresponding canonical correlation  $\mathbf{Z}$  is above a significance threshold as described below.

The significance threshold on the canonical correlations is set by sampling a bootstrap distribution of the maximum canonical correlation using  $\hat{\mathbf{E}}_s$ , the subject-level observation noise identified previously, instead of the subject-level components of interest  $\hat{\mathbf{P}}_s$ .

<sup>2</sup> The singular value spectrum of  $\hat{\mathbf{P}}_s$  is made only of ones:  $\hat{\mathbf{P}}_s \hat{\mathbf{P}}_s^T = \mathbf{1}$ . In other words, the amount of variance each IC which accounts for at the subject level is not modeled at the group level: only the patterns are retained,  $\hat{\mathbf{E}}_s$  is not considered.

Selected canonical variables have a probability  $p < 0.05$  of being generated by data consisting only of observation noise. This procedure corresponds thus to keeping only the group-level components more reproducible than observation noise. We give more details on this procedure in the [supplementary materials](#).

The estimation procedure for the group-level subspace of interest is thus done by minimizing the amount of unexplained signal in each subject while using a fixed number of components, and then by maximizing the subspace stability at the group level. Both optimizations are performed using SVD. For each step, the unexplained variance is chosen based on a noise model.

#### Group-level independent components

The above estimation procedure selects group-level components  $\mathbf{B}$  spanning the subspace of common patterns of activation. We apply FastICA (Hyvärinen and Oja, 2000) on this subspace, to separate group-level independent sources  $\mathbf{A}$ , estimating the mixing model 1.

Finally, as the interesting features of the ICA patterns lie in blobs standing out from the background, it is important to threshold the resulting map to keep only the tail of the intensity distribution.<sup>3</sup> For this, we use as simple null-hypothesis distribution a normal distribution of unit standard deviation. Indeed, the FastICA algorithm works on whitened data (thus of isotropic unit variance) and estimates maximally non-Gaussian directions. As in Schwartzman et al. (2009), this null hypothesis models the central mode of the maps, and we select the voxels for which the absolute intensity exceeds a fixed threshold, that corresponds to a certain level of specificity with respect to null distribution. More complex thresholding approaches have been developed, for instance based on mixture modeling (Beckmann and Smith, 2004). Our method leads to fewer selected voxels on a pattern with very few salient features, such as on artifact patterns. In addition, it is consistent with the FastICA model.

The main differences of our estimation method compared to the most currently used fMRI ICA methods (GIFT and MELODIC) are:

1. Model-order determination by rejecting components that can be generated by Gaussian noise.
2. Selecting a subspace of reproducible components at the group level using canonical correlation analysis as a device to separate subject-level observation noise for group variability.
3. Thresholding ICs based on the absolute value of voxel intensity.

In this article, we are mostly interested in discussing the importance of point 2, which is the expression of the group model used to describe between-subject variability.

#### Cross-validation of group-level patterns

The use of ICA is motivated by the fact that the patterns extracted from the fMRI data display meaningful features in relation to our knowledge of functional neuroanatomy. The validation criteria for an ICA decomposition are unclear, as this algorithm is not based on a testable hypothesis. However, for the method to be usable in group studies, the features extracted from a group of healthy controls should be comparable between different subgroups of subjects, and generalize to different subgroups.

To test the reproducibility of the results across groups of subjects, we split our group of subjects in two and learn ICA maps from each sub-group: this yields the sets of patterns  $\mathbf{A}_1$  and  $\mathbf{A}_2$ . We compare the overlap of thresholded maps and reorder one set to match maps by maximum overlap. Reproducibility can be quantified by studying the cross-correlation matrix  $\mathbf{C} = \mathbf{A}_1^T \mathbf{A}_2$ . For unit-normed components,  $\mathbf{C}_{ij} = 1$  if and only if  $(\mathbf{A}_1)_i$  and  $(\mathbf{A}_2)_j$  are identical.

<sup>3</sup> Although this is sometimes done implicitly by simply referring to the brain regions that stand out when looking at spatial maps, it is important to this do this reduction explicitly.

To quantify the reliability of the patterns identified on the full datasets, we select for each pattern extracted from the full dataset the best matching one in the different subsets computed in the cross-validation procedure using Pearson's correlation coefficient. Along with the extracted maps, we report the average value of this pattern-reproducibility measure.

#### Validation experiments: comparing to other methods

We compare our method with state-of-the-art methods using different group-model estimation procedures: concatenation approaches using the implementations of the GIFT ICA toolbox (Calhoun et al., 2001) and the MELODIC software (Smith et al., 2004), as well as the tensor ICA model, also implemented in the MELODIC software (Beckmann and Smith, 2005).

Also, to isolate the effect of separating subject-level variance from group-level variance in the estimation algorithm from implementation-specific details, we run a modified version of CanICA without the CCA step, retaining the subject-level variance  $\Sigma_s$  at the group level on the same datasets: Eqs. (15) and (16) are thus replaced by

$$\tilde{\mathbf{P}}_s = (\Sigma_s \mathbf{V}_s)_{1 \dots n_{\text{subj}}}, \quad (19)$$

$$\tilde{\mathbf{W}}_s = (\mathbf{U}_s)_{1 \dots n_{\text{subj}}}. \quad (20)$$

With this estimation algorithm, the variance explained by the components at the subject level is carried on to the group level. This is analogous to a fixed-effect model.

As estimating larger numbers of independent components tend to explore less stable patterns, we run all analysis using the number of components estimated by CanICA. We perform cross-correlation analysis on the non-thresholded ICA maps, but also use each implementation's thresholding algorithm to separate the features of interest. In an effort to separate the impact of the group model from the impact of the thresholding heuristic, we study reproducibility of non-thresholded maps, as well as thresholded maps, although non-thresholded maps can be considered as of little neuro-scientific interest. The mixture model implemented in MELODIC selects 3000 voxels on average on all the patterns. We set the specificity for the other methods to select the same average number of voxels. For GIFT, the thresholding is done on the  $t$ -statistics maps generated by the algorithm. We threshold at  $|t| > 2$ .

#### Similarity measures for decompositions learned in different subjects

We define two measures to quantify the stability of the subspace spanned by the components and the reproducibility of the maps. First, a measure of the overlap of the subspaces selected in both groups is given by the energy of the matrix, i.e. its Frobenius norm:  $E = \text{tr}(\mathbf{C}^T \mathbf{C})$ . To compare this quantity for different subspace sizes, we normalize it by the minimum dimension of the subspaces,

$$d = \min(\text{rank } \mathbf{A}_1, \text{rank } \mathbf{A}_2) \quad (21)$$

$$e = \frac{1}{d} \text{tr}(\mathbf{C}^T \mathbf{C}). \quad (22)$$

$e$  quantifies the reproducibility of the subspace spanned by the maps. For  $e = 1$ , the two groups of maps span the same subspace, although individual independent components may differ. If  $\mathbf{A}_1$  and  $\mathbf{A}_2$  have different dimensions, taking the minimum dimension of both does not account for possible instability of patterns extracted from a group and not the other. However, the dimensionality as estimated by our procedure does not vary much over the cross-validation pairs (maximum 10% relative difference).

Second, we measure one-to-one reproducibility of maps. We reorder the matrix  $\mathbf{C}$  by matching sequentially maximally mutually correlated components of  $\mathbf{A}_1$  and  $\mathbf{A}_2$ , considering the absolute value of

the pair-wise correlation. This procedure creates a reordered cross-correlation matrix  $\tilde{\mathbf{C}}$  with maximal matching values on the diagonal. If one of the two groups has more components than the other, we use the smaller of the two as a reference, and stop the matching procedure once all of its components have been matching. As a result,  $\tilde{\mathbf{C}}$  is a square matrix that we populate with absolute values of pair-wise correlations. Although we acknowledge that this algorithm may not give the optimal matching in general, we observe that the solution is satisfactory in practice. We use the normalized trace of the reordered cross-correlation matrix,

$$t = \frac{1}{d} \text{tr}(\tilde{\mathbf{C}}), \quad (23)$$

as an overall measure of one-to-one overlap between matched pairs of components.

#### fMRI datasets

We apply the probabilistic ICA method described above to extract brain networks from two multi-subjects datasets: a group resting-state study and a functional localizer.

#### Resting-state data

We use datasets from a resting-state experiment: subjects were blindfolded and instructed to keep their eyes closed. The resting state session was recorded as the first session preceding a series of cognitive experiments as part of a study not detailed here. At the time of resting-state data collection, subjects were naive with respect to the nature of the subsequent experiments. Twelve right-handed healthy volunteers (two females; ages 19–30) gave written informed consent before participation in imaging on a 3 T MRI whole-body scanner (Tim-Trio, Siemens, Erlangen). The study received local ethics committee approval. 820 EPI volumes (25 slices,  $TR = 1.5$  s,  $TE = 30$  ms, FOV 19.2 cm  $\times$  19.2 cm, 3 mm isotropic resolution) were acquired during a rest period of 20 min. A full description of the paradigm and the acquisition parameters can be found in Sadaghiani et al. (2009).

#### Functional localizer data

We used an event-related experimental paradigm consisting of 10 conditions. Subjects were presented with a series of stimuli and were engaged in tasks such as passive viewing of horizontal or vertical checkerboards, left or right button press after auditory or visual instruction, computation (subtraction) after video or visual instruction and sentence processing, from the auditory or visual modality. Events were randomly occurring in time (mean inter stimulus interval: 3 s), with 10 occurrences per event type (except button presses for which there are only five trials per session).

Two hundred right-handed subjects participated in the study. The subjects gave informed consent and the protocol was approved by the local ethics committee. Functional images were acquired on a 3 T Bruker scanner using an EPI sequence (40 slices,  $TR = 2.4$  s,  $TE = 60$  ms, FOV 19.2 cm  $\times$  19.2 cm, 3 mm isotropic resolution). A session comprised 132 scans. The first four functional scans were discarded to allow the MR signal to reach steady state. The experimental paradigm and the acquisition parameters are described with more detail in Pinel et al. (2007). To balance the comparison with resting-state data, we first report reproducibility analysis on 12 subjects acquired consecutively. We also study reproducibility on a group of 62 subjects all acquired successively in time.

#### Preprocessing

The above datasets were preprocessed using the SPM5 software (Wellcome Department of Cognitive Neurology; <http://www.fil.ion.ucl.ac.uk/spm>). After slice-timing interpolation and motion correction, cerebral volumes were realigned to the MNI152 inter-subject

template and smoothed with a 5 mm isotropic Gaussian kernel. Voxels contained in a mask of the brain (approximately 40 000 voxels) were selected for pattern extraction using the CanICA procedure. We do not apply a frequency filter to the time-series, as we have found out that this yields ICs displaying better separation between brain networks and artifacts such as movement or blood-flow related patterns.

## Results

### Group-level patterns extracted

#### Resting-state dataset

On the resting-state dataset, consisting of 820 scans, CanICA identified in average 50 non-observation-noise principal components at the subject level and a subspace of 42 reproducible patterns at the group level (see Eq. (7)). This number matches those commonly hand-selected by users in current ICA studies. Kiviniemi et al. (2009) reported 60 stable independent components and 42 brain networks in a study of group-level ICA patterns stability on a similar dataset. On this dataset made of long time sequences, model-evidence-based methods such as those used in Beckmann and Smith (2004) select 200 components at the subject level (or 340 when applied on smoothed data). Out of the 42 maps extracted by the CanICA, we identified by eye 26 putatively components as brain networks (Fig. 2). When using CanICA without CCA, we putatively identified only 11 components extracted as brain networks. Other components related to physiological noise or movement form patterns in the BOLD signal common to the group as they are related to reproducible anatomical features.

#### Functional localizer dataset

On the functional localizer dataset, consisting of 150 scans, CanICA extracts 20 ICA patterns, out of which we identify 13 putative functional networks (see Fig. 3). When using Eq. (20) instead of Eq. (15), i.e. without the whitening of the individual patterns imposed in the CCA, only 6 functional networks are identified (see supplementary materials).

### Reproducibility results

The reproducibility metrics we obtained by cross-validation are reported in Table 1. For the resting-state experiment, on non-thresholded maps, the GIFT and ConcatICA perform similarly with regards to subspace-stability, which can be explained by the fact that they implement similar group models. For this experiment, the reproducibility of the TensorICA model is not as good. Conversely, for the localizer experiment, the reproducibility of TensorICA is good. The assumption of the TensorICA model that networks share similar time courses across subjects is clearly more suitable for task-driven studies. The CCA-based estimation procedure of the two-level group model of CanICA yields the most stable subspace in both experiments. However, the relative performance of different methods changes between the two experiments that have very different signal length and TR.

Performance of the method with regards to one-to-one reproducibility of thresholded maps does not differ significantly from non-thresholded maps for all the methods using histogram-based thresholding. On the other hand, GIFT thresholds a  $t$ -statistic map over back-reconstructed subject-specific components, which is quite unstable across population splits. As a result, one-to-one matching of thresholded maps extracted by GIFT does not perform well. We also note that the independent components extracted from the functional localizer dataset are less contrasted than those extracted from the resting-state data because the number of volumes is smaller, and as a result thresholding has a more detrimental impact on stability.

The distribution of maximal component matching from one subgroup to another provides an assessment of the reproducibility of the individual maps. Table 2 gives the percentile of this distribution on both datasets, for thresholded and non-thresholded maps, when these are extracted with CanICA. The maps for which there is a matching above 50% are of particular interest, as they have sufficient correspondence to establish a one-to-one mapping with a simple matching scheme.

In addition, we have performed the same reproducibility analysis using a higher number of independent components, to compare the different methods for parameters similar to the analysis of Kiviniemi et al. (2009) (see Table 4). As expected, the reproducibility of the selected subspace is reduced compared to using a small number of components selected, but the relative performance of the methods is similar.

Finally, we have studied reproducibility for a large group of subjects, on the localizer dataset for CanICA with and without CCA. As can be seen in Table 3, reproducibility is improved when estimating ICs on larger groups. On these groups, the use of CCA remains an important factor of reproducibility.

We conclude from this cross-validation study that CanICA is a suitable tool for purely data-driven extraction of stable markers from fMRI data on a homogeneous group of subjects as it yields a small number of highly repeatable features that can be identified between different groups. The procedure is fully automatic as it does not rely on identification of previously known activation patterns corresponding to cognitive paradigms. Unlike previous reproducibility studies (e.g. Damoiseaux et al., 2006), we report on the complete set of patterns extracted.

## Discussion

### Factors impacting reproducibility

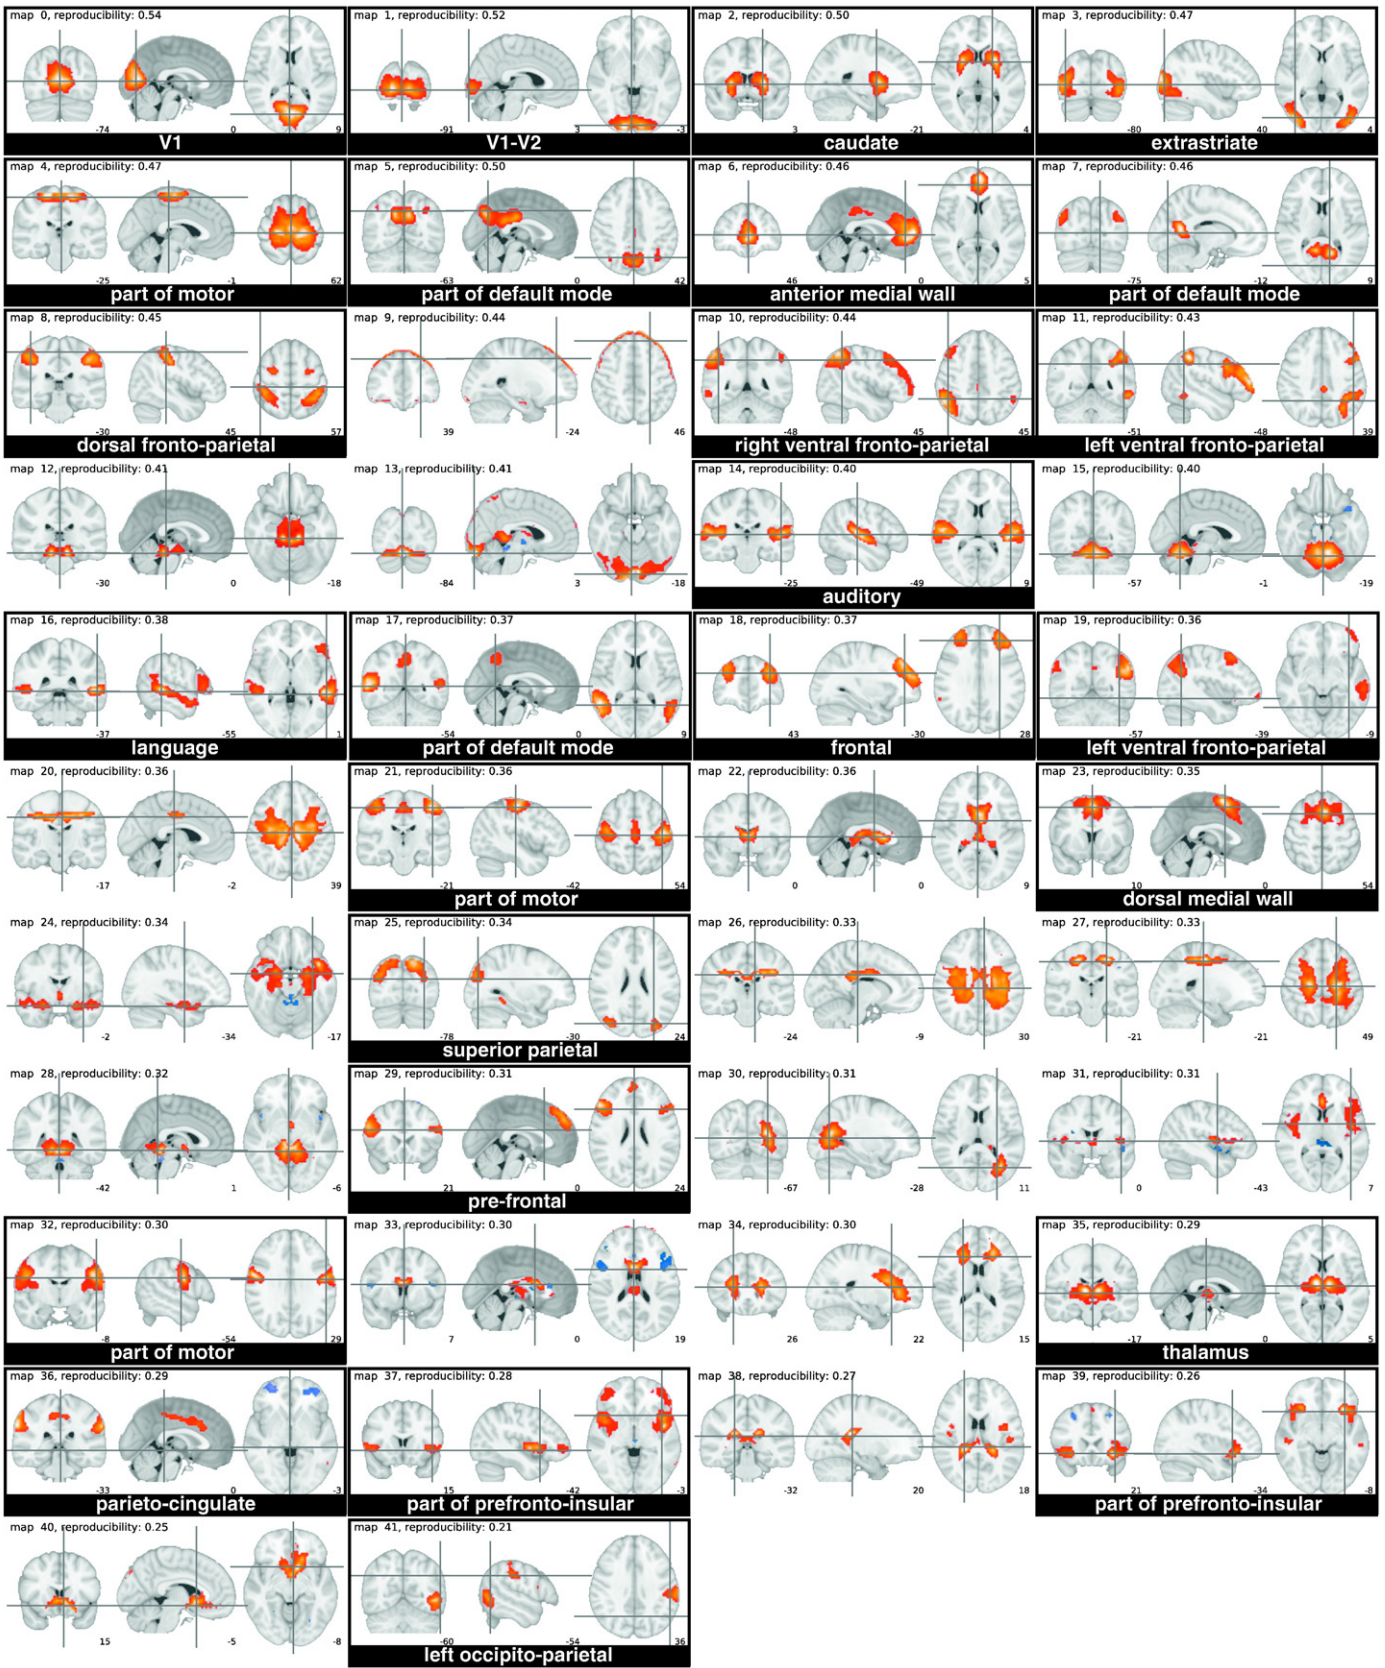
#### Importance of the signal subspace

We have used two different metrics to measure reproducibility of the ICs. While  $t$  measures exact one-to-one matching of the final ICs,  $e$  is a measure of reproducibility of the subspace, and is thus independent of the ICA step. The FastICA algorithm was ran with different parameters (in symmetric or in deflation mode, and using either the cube or the logcosh non-linearity), and gave similar reproducibility results as measured by  $t$  on both datasets ( $t$  metrics differing of .01). While GIFT and MELODIC in TensorICA use different ICA algorithms, MELODIC in concatenation mode and CanICA, with and without CCA, all use the FastICA algorithm. We thus conjecture that the difference between the reproducibility scores of these last methods is mainly due to different subspace-selection procedures which consist of preprocessing and group-model estimation.

#### Importance of thresholding on canonical correlations

We compared four different sub-space selection procedures: CanICA with and without CCA, MELODIC and GIFT. CanICA without CCA implements a fixed-effect model, in which the principal components of each subjects are concatenated without whitening before group-level analysis. CanICA with CCA uses whitening of the individual datasets to perform canonical correlation analysis and select group-level components via a well-known reproducibility score: the canonical correlation. The comparison of performing analysis with and without CCA over different datasets (Table 1), with varying group sizes (Table 3), and for different number of ICs (Table 4), shows the importance of the CCA step.

To our knowledge, there is no published detailed description of the group-level estimation procedure implemented in MELODIC ConcatICA and GIFT, we therefore base our discussion on our analysis of the software packages. GIFT groups individual subject datasets and performs successive PCAs and thresholding. This can be interpreted



**Fig. 2.** The 42 ICA maps extracted by CanICA from the resting-state dataset (display in radiologic convention: the right hemisphere is on the left of the axial view). The maps are ordered by reproducibility (from left to right on each line, and top to bottom line after line). Maps corresponding to functionally plausible networks are in a black frame, whereas maps likely corresponding to artifacts are not framed. Extracted brain networks are labeled with the name of the general structure they can be related to.



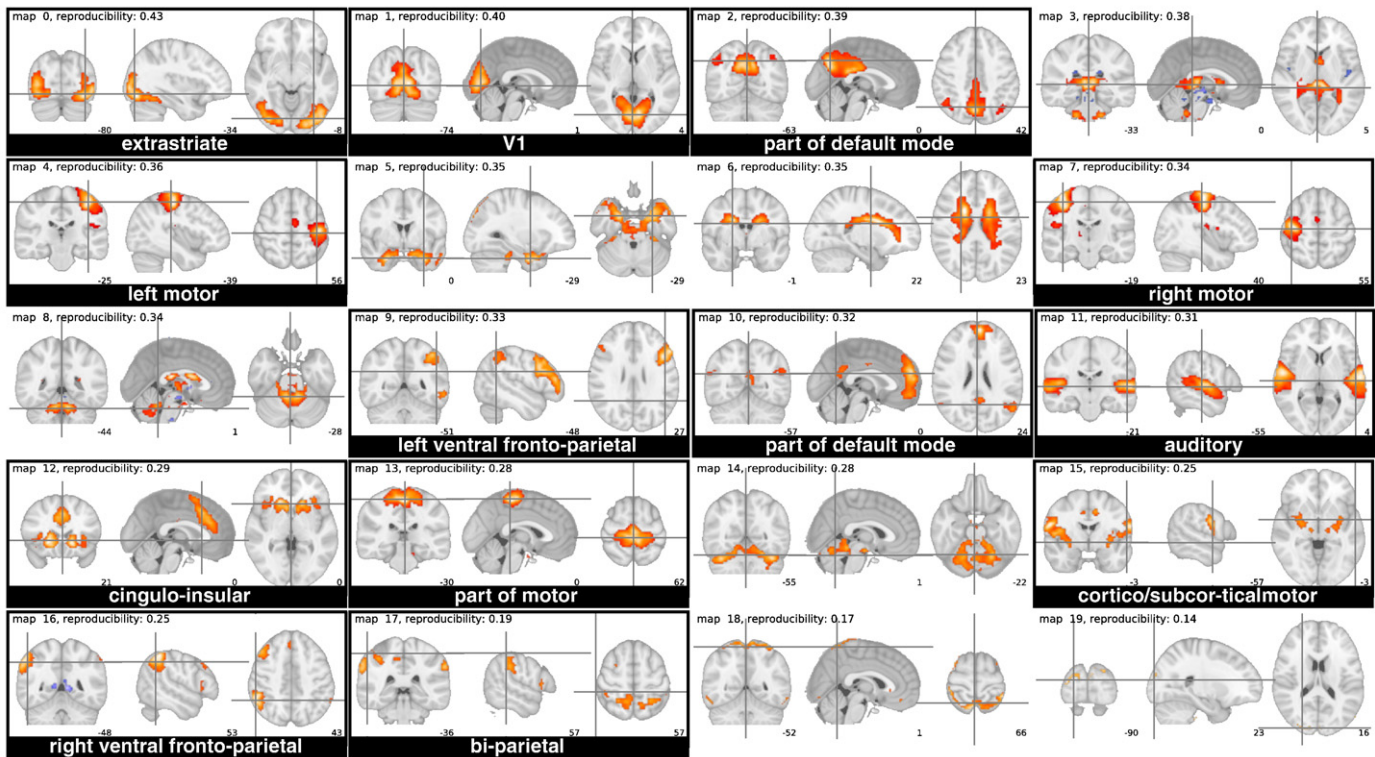


Fig. 3. The 20 ICA maps extracted by CanICA on the functional localizer dataset (radiologic convention).

as applying nested fixed-effects model. MELODIC applies a group-average filtering matrix before performing whitening and group-level data reduction via SVD. For these methods, the group-level components are thus not selected using canonical correlation analysis on the individual subject data. While canonical correlation selects the sub-space to optimize for reproducibility, the procedures for group-level data reduction applied by both MELODIC and GIFT impose some reproducibility criteria, as can be seen by their  $e$  score on the resting-state dataset which out-performs a simple concatenation (see Table 1).

We note that when performing analysis with a large number of ICs (Table 4), selecting the signal subspace on a canonical correlation criteria is less critical, as the retained subspace covers a larger proportion of the initial signal. On the opposite, on the localizer dataset, performing CCA is important for reproducibility even for large groups of subjects (Table 3).

Table 1

Average reproducibility measures  $e$  and  $t$  for Group ICA, Tensor ICA and CanICA calculated on the half-split cross-correlation matrices, both for non-thresholded and thresholded maps. Numbers in parenthesis give the standard deviation across the different splits.

	MELODIC TensorICA	GIFT GroupICA	MELODIC ConcatICA	CanICA no CCA	CanICA CCA
<i>Resting state</i>					
Unthresholded ICA maps					
$e$ : subspace stability	.47 (.06)	.58 (.04)	.58 (.04)	.36 (.02)	.71 (.01)
$t$ : one-to-one matching	.36 (.03)	.53 (.04)	.51 (.04)	.36 (.02)	.72 (.05)
Thresholded ICA maps					
$t$ : one-to-one matching	.35 (.02)	.10 (.01)	.50 (.03)	.21 (.02)	.62 (.04)
<i>Localizer</i>					
Unthresholded ICA maps					
$e$ : subspace stability	.54 (.05)	.25 (.03)	.43 (.02)	.35 (.01)	.52 (.01)
$t$ : one-to-one matching	.36 (.02)	.34 (.04)	.35 (.03)	.37 (.02)	.55 (.02)
Thresholded ICA maps					
$t$ : one-to-one matching	.29 (.03)	.02 (.01)	.31 (.05)	.26 (.03)	.46 (.02)

### Thresholding heuristic

The metric  $t$  applied to thresholded ICs measures the reproducibility of features identified by the thresholding heuristic. This measure indicates the ability of the ICA algorithm to yield reproducible salient features, but is also a factor of the thresholding heuristic. The focus of this paper is not to discuss thresholding heuristics, and we use a simple one. However, we note that the impact of different heuristics (a mixture model, as in MELODIC, or amplitude-based thresholding, as in CanICA) on the cross-validation reproducibility score varies across the different studies. In Table 1, mixture models seem to impact reproducibility to a lesser extent, whereas when using a high number of ICs (Table 4), amplitude-based thresholding performs better. This can be understood by the fact that the heuristics depend strongly on the histogram of the maps, which vary with the number of ICs or subjects. It is thus difficult to conclude on thresholding without further study.

### Interpretation of the ICA patterns

#### Sensitivity of the method to brain networks

The patterns extracted display many different components that can be interpreted as functional networks. On the resting-state data, this is the case for 26 ICs out of 42 (Fig. 2). MELODIC's Concat-ICA, which corresponds to a state-of-the-art implementation of a one-level

Table 2

Percentiles of maximum one-to-one matching between maps extracted with CanICA on two sub-groups of six different subjects.

	Unthresholded		Thresholded	
	RS	Loc.	RS	Loc.
Matching above 50%	91%	64%	77%	44%
Matching below 25%	0%	2%	5%	25%
Matching above 75%	55%	19%	34%	28%

**Table 3**  
Reproducibility scores with using CanICA on the localizer dataset with many subjects.

	12 subjects		20 subjects		40 subjects		62 subjects	
	No CCA	CCA	No CCA	CCA	No CCA	CCA	No CCA	CCA
<i>Unthresholded ICA maps</i>								
<i>e</i> : subspace stability	.35 (.01)	.52 (.01)	.36 (.01)	.57 (.01)	.42 (.01)	.71 (.01)	.50 (.01)	.78 (.01)
<i>t</i> : one-to-one matching	.37 (.02)	.55 (.02)	.40 (.02)	.57 (.03)	.45 (.01)	.68 (.03)	.49 (.01)	.72 (.03)
<i>Thresholded ICA maps</i>								
<i>t</i> : one-to-one matching	.26 (.03)	.46 (.02)	.29 (.02)	.45 (.03)	.33 (.02)	.56 (.03)	.38 (.03)	.60 (.04)

group model, when run on the same dataset with the same model order, yields 20 identifiable functional networks (Fig. 5, supplementary materials). Using the same implementation as CanICA, with the same preprocessing and thresholding steps, but without modeling two levels of variance (fixed-effect model), only 11 patterns extracted can be identified as brain networks.

However, the brain networks detected in both the fixed-effects and the random-effects group models display close resemblance: modeling two levels of variance, at the subject and at the group level, does not significantly alter the group-level brain networks extracted. The difference in extracted networks corresponds to a difference in sensitivity of the methods to brain networks. This can be understood by the fact that different networks can be activated in varying proportion in each subject. This variability is tamed by the whitening introduced in the CCA step that gives an equal weighting to all the information captured in each dataset. As a result more networks are identified: well-known networks can be segmented into reproducible sub-networks, and additional networks seldom encountered in ICA analysis are extracted.

We note that on the resting-state dataset, because of the reduced field of view, the upper part of the cortex is cropped (as can be seen on map 4), and we have no information on the networks in the lower temporal lobes and in the cerebellum.

#### Extracted brain networks

Most of the networks extracted can be related to stable networks well-known from the resting-state literature, such as the visual system or the fronto-parietal ventral and dorsal structures labeled as attentional networks (Fox et al., 2006). However, due to the increased statistical power, we resolve sub-structures of these networks.

For instance, the network known as *default mode network* appears to be split into several sub-networks separating portions of the posterior cingulate region, as well as occipito-parietal junction, precuneus and medial pre-frontal cortex (maps 5, 6, 7, and 18 on

the resting-state dataset and 2 and 10 on the localizer dataset). In particular, the retrosplenial cortex stands separated from the posterior cingulate cortex, grouped on map 7 of the resting-state dataset with parietal regions. It has been shown to be the focus of anatomical connections to the medial pre-frontal regions (Greicius et al., 2009). Also the ventral anterior cingulate cortex, shown on map 6 of the resting-state dataset, has been consistently identified separately from the default mode network, with evidence from EEG measurements that it forms a differentiable network (Mantini et al., 2007). This sub-network has been associated with self-referential mental activity (D'Argembeau et al., 2005; Johnson et al., 2006).

The networks identified in the frontal lobes, associated to executive functions, display considerable variability in the literature as well as between our datasets. Maps 37 and 39 on the resting-state dataset and map 12 on the functional localizer correspond to the network related to salience processing in Seeley et al. (2007) and associated with task set maintenance in Dosenbach et al. (2006, 2007, 2008). As reported by Dosenbach et al. (2007) and Seeley et al. (2007), it forms a distinct network from the parietal-frontal network made of maps 10, 19 and 29 of the resting-state dataset.

Finally, we extract from the resting-state dataset some function-specific networks such as the putative language network (map 16) or an occipito-parietal network (map 41) related to the dorsal visual pathway. Map 15 of the localizer dataset is a rich, anatomically well-defined cortico-subcortical motor network that comprises the ventral and angular motor cortices, as well as the insular cortex and the lentiform nucleus and thalamus.

#### Occurrence of the brain networks across groups

Different networks can be recruited depending on the task, and thus we may expect the occurrence of cognitive networks to vary when estimated on different experiments, or different sessions. For instance, we observe that the right and left motor cortices appear on different independent components for the localizer experiment, which may be explained by the separate right and left finger-tapping tasks present in the experimental paradigm of the localizer experiment, whereas in the resting-state experiment, the motor areas are divided symmetrically in somatotopic regions corresponding to lower body, upper body and face.

The reproducibility number, indicated in Figs. 2 and 3, gives an indication of the occurrence of a network identified across subgroups. Visual and attentional areas stand out among the most reproducible brain networks in both datasets—different parts of the visual cortex, and the network coined *visuo-spatial* system in Beckmann et al. (2005). Parietal regions are embedded into various networks, bilateral or not, and possibly but not necessarily associated with frontal regions, according to the dataset considered or the processing strategy. The same kind of observation also applies to frontal regions. Bilateral auditory cortices, on the other hand, appear as a robust single network across datasets. In general, networks related to primary areas (sensory or motor) are more stable across groups, paradigms, and methods, whereas networks related to higher-level areas (mainly in the frontal lobes) or higher-level specific tasks (such as language

**Table 4**  
Average reproducibility measures *e* and *t* for Group ICA, Tensor ICA and CanICA calculated on the half-split cross-correlation matrices, with 70 ICs, both for non-thresholded and thresholded maps. Numbers in parenthesis give the standard deviation across the different splits. The localizer dataset is made of runs of 132 volumes, selecting 70 ICs explores thus the tail of the PCA.

	MELODIC TensorICA	GIFT GroupICA	MELODIC ConcatICA	CanICA No CCA	CanICA CCA
<i>Resting state</i>					
Unthresholded ICA maps					
<i>e</i> : subspace stability	.28 (.01)	.41 (.01)	.53 (.02)	.28 (.01)	.49 (.02)
<i>t</i> : one-to-one matching	.15 (.01)	.37 (.01)	.47 (.02)	.28 (.01)	.50 (.01)
Thresholded ICA maps					
<i>t</i> : one-to-one matching	.19 (.02)	.06 (.01)	.33 (.02)	.29 (.01)	.57 (.02)
<i>Localizer</i>					
Unthresholded ICA maps					
<i>e</i> : subspace stability	.52 (.07)	.17 (.02)	.41 (.01)	.25 (.01)	.30 (.01)
<i>t</i> : one-to-one matching	.28 (.01)	.21 (.02)	.29 (.03)	.27 (.02)	.34 (.02)
Thresholded ICA maps					
<i>t</i> : one-to-one matching	.23 (.02)	.02 (.01)	.28 (.03)	.35 (.03)	.39 (.01)

processing) are less stable, and sometimes ill-resolved by the methods.

Although the comparison of resting-state versus activation datasets is not the topic of this paper, one can notice that the results are relatively consistent between the two datasets, and that the resting-state data based on longer time series tends to produce a larger number of interpretable components. This hints at the fact that these networks may be defined not only at rest, but also in much more general contexts of co-activation (Smith et al., 2009).

#### Limitations of the method

When running a group analysis using CanICA, signal originating from different subject are put in correspondence via spatial alignment (performed by the preprocessing steps). While the CanICA model accounts for subject-to-subject differences, one of its major limitations is that it does not model spatial variability across subjects. This is why the estimation is applied on smoothed data. Similarly, the validation metrics  $e$  and  $t$  only test for spatial correspondence. Differences between two ICs estimated on different subgroups that can be accounted for by a spatial displacement of features will induce poor reproducibility scores, and will hinder the matching procedure used in the  $t$  reproducibility score.

A fundamental limitation of the ICA-based method presented and the validation metrics is that they make no difference between neuronal and artifactual signal. Thus, not only will the method extract and report non-neuronal signal, but the corresponding maps will impact the validation metrics. Additional pattern-matching techniques can be used for these purposes (De Martino et al., 2007). Extracting markers of the non-neuronal signal can also have some value, for instance to use them as nuisance regressors (Perlberg et al., 2007).

Choosing the model order of an ICA method is a difficult problem, as the ICA algorithm will estimate orthogonal components spanning all the signal subspace with no measure of statistical relevance. Our approach tries to combine criteria of non-Gaussianity at the subject level, and a reproducibility threshold on the signal subspace at the group level to identify the reproducible non-Gaussian signal subspace. One limitation of selecting a subspace of reproducible signal across subjects is that projection on this subspace may lead to projecting on the same IC components that could be separated in larger subspace, as outlined by Kiviniemi et al. (2009). This can be seen via component splitting when using higher model orders. On the other hand, selecting a higher number of components leads to less reproducible components and uncontrolled ICs that are considered as artifacts and not interpreted in most analysis.

In the CCA estimation step, the statistical measure of reproducibility (canonical correlation) is a linear correlation measure and is not informed with regards to the criteria of ICA, non-Gaussianity. This is a limitation of the canonical correlation approach in our framework. One could consider kernel or non-linear CCA models that would be more specific to ICA criteria. Such procedures are much less tractable, and there are less known results in statistics.

Finally, while CanICA models group-variability during the estimation step to extract ICs representative of the group, it does not provide a procedure to infer individual components related the group-level maps. For this purpose, we suggest using dual regression as in Filippini et al. (2009): the dual regression framework is independent of the estimation procedure used to extract maps representative of the group.

#### Algorithm complexity and numerical efficiency

Because the estimation of the group-level model relies solely on simple linear algebra routines, and the ICA optimization loop is performed on a small number of selected components, it can be very efficient on large data when implemented with optimized linear

algebra packs, both in terms of number of operations and memory. The computational cost of estimating the full model on a group of  $S$  subjects, with  $m$  volumes each, and  $p$  voxels in the brain is dominated by the cost of the subject-level PCA that scales in  $S \cdot m^2 \cdot p$ . The group-level inference is made of the CCA step, that scales in  $S^2 \cdot p$ , and the ICA step, that scales in  $S \cdot p$ . For our data set, it takes a few minutes on a 2 GHz Intel core Duo. The speed of this step is critical for cross-validation, as the initial subject-level model does not need to be recomputed.

Performance is important to scale to long fMRI time series, high-resolution data, or large groups. In addition, as the group-level pattern extraction (CCA and ICA) is very fast, cross-validation of the group patterns is feasible on modest hardware.

On the other hand, the model-order selection steps imply bootstrap analysis. These steps are computationally costly. In particular estimation of the number of principal components retained at the subject level involves evaluation of many SVDs, a lengthy process (scaling in  $p \cdot m^2$ ). We do not estimate the subject-level model order for each sub-group during the cross-validation of the group patterns.

#### Conclusion

We have presented a multivariate two-level generative model for multi-subject datasets and applied it to an ICA model and corresponding pattern-extraction algorithm for fMRI data, CanICA. Compared to existing methods, our approach uses non-parametric noise description for model-order selection. As a result, the method is auto-calibrated and extracts in a fully automated way meaningful and reproducible features from fMRI data. In addition, we have introduced a cross-validation procedure and associated metrics for ICA patterns and used it to establish validity of group-level maps.

ICA is an unstable procedure with no intrinsic significance testing, but we have shown that our pattern-extraction method, based on a mixed-effects-like group model, can yield a set of thresholded maps of which many are reproducible and can be identified one to one when extracted from two different groups of only six healthy controls. Reproducibility is an important feature of exploratory analysis methods, as the validity of their results cannot be established by hypothesis testing. In addition, group reproducibility on control groups and one-to-one matching between groups is necessary when using extracted patterns as bio-markers for group analysis.

#### Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at [doi:10.1016/j.neuroimage.2010.02.010](https://doi.org/10.1016/j.neuroimage.2010.02.010).

#### References

- Beckmann, C.F., Smith, S.M., 2004. Probabilistic independent component analysis for functional magnetic resonance imaging. *Trans. Med. Imaging* 23 (2), 137–152.
- Beckmann, C.F., Smith, S.M., 2005. Tensorial extensions of independent component analysis for multisubject FMRI analysis. *Neuroimage* 25 (1), 294–311 Mar.
- Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360 (1457), 1001–1013 May.
- Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34 (4).
- Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10, 186–198.
- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001. A method for making group inferences from functional MRI data using independent component analysis. *Hum. Brain Mapp.* 14 (3), 140–151 Nov.
- Cordes, D., Haughton, V.M., Arfanakis, K., Wendt, G.J., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. *Am. J. Neuroradiol.* 21 (9), 1636–1644.

- Cordes, D., Haughton, V., Carew, J.D., Arfanakis, K., Maravilla, K., 2002. Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magn. Reson. Imaging* 20 (4), 305–317.
- Damoiseaux, J.S., Rombouts, S.A.R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. *Proc. Natl. Acad. Sci. U. S. A.* 103 (37), 13848–13853.
- D'Argembeau, A., Collette, F., Van der Linden, M., Laureys, S., Del Fiore, G., Degueldre, C., Luxen, A., Salmon, E., 2005. Self-referential reflective activity and its relationship with rest: a PET study. *Neuroimage* 25 (2), 616–624.
- Daubechies, I., Roussos, E., Takerkart, S., Benharrosh, M., Golden, C., D'Ardenne, K., Richter, W., Cohen, J.D., Haxby, J., 2009. Independent component analysis for brain fMRI does not select for independence. *Proc. Natl. Acad. Sci. U. S. A.* 106 (26), 10415–10422.
- De Martino, F., Gentile, F., Esposito, F., Balsi, M., Di Salle, F., Goebel, R., Formisano, E., 2007. Classification of fMRI independent components using IC-fingerprints and support vector machine classifiers. *Neuroimage* 34 (1), 177–194.
- Dosenbach, N.U.F., Visscher, K.M., Palmer, E.D., Miezin, F.M., Wenger, K.K., Kang, H.C., Burgund, E.D., Grimes, A.L., Schlaggar, B.L., Petersen, S.E., 2006. A core system for the implementation of task sets. *Neuron* 50 (5), 799–812.
- Dosenbach, N.U.F., Fair, D.A., Miezin, F.M., Cohen, A.L., Wenger, K.K., Dosenbach, R.A.T., Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., et al., 2007. Distinct brain networks for adaptive and stable task control in humans. *Proc. Natl. Acad. Sci. U. S. A.* 104 (26), 11073.
- Dosenbach, N.U.F., Fair, D.A., Cohen, A.L., Schlaggar, B.L., Petersen, S.E., 2008. A dual-network architecture of top-down control. *Trends Cogn. Sci.* 12 (3), 99–105.
- Esposito, F., Scarabino, T., Hyvarinen, A., Himberg, J., Formisano, E., Comani, S., Tedeschi, G., Goebel, R., Seifritz, E., Di Salle, F., 2005. Independent component analysis of fMRI group studies by self-organizing clustering. *Neuroimage* 25 (1), 193–205.
- Filippini, N., MacIntosh, B.J., Hough, M.G., Goodwin, G.M., Frisoni, G.B., Smith, S.M., Matthews, P.M., Beckmann, C.F., Mackay, C.E., 2009. Distinct patterns of brain activity in young carriers of the APOE-ε4 allele. *Proc. Natl. Acad. Sci. U. S. A.* 106 (17), 7209.
- Fox, M., Raichle, M., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* 8, 700–711.
- Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci. U. S. A.* 102 (27), 9673–9678.
- Fox, M.D., Corbetta, M., Snyder, A.Z., Vincent, J.L., Raichle, M.E., 2006. Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proc. Natl. Acad. Sci. U. S. A.* 103 (26), 10046.
- Friston, K.J., Stephan, K.E., Lund, T.E., Morcom, A., Kiebel, S., 2005. Mixed-effects and fMRI studies. *Neuroimage* 24 (1), 244–252.
- Garrity, A.G., Pearlson, G.D., McKiernan, K., Lloyd, D., Kiehl, K.A., Calhoun, V.D., 2007. Aberrant “default mode” functional connectivity in schizophrenia. *Am. J. Psychiatry* 164 (3), 450.
- Golland, P., Golland, Y., Malach, R., 2007. Detection of spatial activation patterns as unsupervised segmentation of fMRI data. *Lect. Notes Comput. Sci.* 4791, 110.
- Greicius, M.D., 2008. Resting-state functional connectivity in neuropsychiatric disorders. *Curr. Opin. Neurol.* 21 (4), 424.
- 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 (1), 253–258.
- Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V., 2004. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *Proc. Natl. Acad. Sci. U. S. A.* 101 (13), 4637–4642.
- Greicius, M.D., Supekar, K., Menon, V., Dougherty, R.F., 2009. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cereb. Cortex* 19, 72–78.
- Guo, Y., Pagnoni, G., 2008. A unified framework for group independent component analysis for multi-subject fMRI data. *Neuroimage* 42 (3), 1078–1093.
- Harshman, R.A., 1970. Foundations of the PARAFAC procedure: models and conditions for an “explanatory” multi-modal factor analysis. *UCLA Work. Pap. Phon.* 16 (1), 1.
- Himberg, J., Hyvarinen, A., Esposito, F., 2004. Validating the independent components of neuroimaging time series via clustering and visualization. *Neuroimage* 22 (3), 1214–1222.
- Honey, C.J., Kötter, R., Breakspear, M., Sporns, O., 2007. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proc. Natl. Acad. Sci. U. S. A.* 104 (24), 10240.
- Hyvarinen, A., Oja, E., 2000. Independent component analysis: algorithms and applications. *Neural Netw.* 13 (4–5), 411–430.
- Johnson, M.K., Raye, C.L., Mitchell, K.J., Touryan, S.R., Greene, E.J., Nolen-Hoeksema, S., 2006. Dissociating medial frontal and posterior cingulate activity during self-reflection. *Soc. Cogn. Affect. Neurosci.* 1 (1), 56.
- Kettenring, J.R., 1971. Canonical analysis of several sets of variables. *Biometrika* 58 (3), 433–451.
- Kiviniemi, V., Kantola, J.H., Jauhiainen, J., Hyvarinen, A., Tervonen, O., 2003. Independent component analysis of nondeterministic fMRI signal sources. *Neuroimage* 19 (2), 253–260.
- Kiviniemi, V., Starck, T., Remes, J., Long, X., Nikkinen, J., Haapea, M., Veijola, J., Moilanen, I., Isohanni, M., Zang, Y.F., et al., 2009. Functional segmentation of the brain cortex using high model order group PCA. *Hum. Brain Mapp.* 47, 194. (Presented at OBHM).
- Krzyszowski, W.J., 1979. Between-groups comparison of principal components. *J. Am. Stat. Assoc.* 74 (367), 703–707 September.
- Langers, D.R.M., 2009. Blind source separation of fMRI data by means of factor analytic transformations. *Neuroimage* 47 (1), 77–87.
- Laufs, H., Kleinschmidt, A., Beyerle, A., Eger, E., Salek-Haddadi, A., Preibisch, C., Krakow, K., 2003. EEG-correlated fMRI of human alpha activity. *Neuroimage* 19 (4), 1463–1476.
- Li, Y.O., Adali, T., Calhoun, V.D., 2007. Estimating the number of independent components for functional magnetic resonance imaging data. *Hum. Brain Mapp.* 28 (11).
- Lowe, M.J., Mock, B.J., Sorenson, J.A., 1998. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *Neuroimage* 7 (2), 119–132.
- Mantini, D., Perrucci, M.G., Del Gratta, C., Romani, G.L., Corbetta, M., 2007. Electrophysiological signatures of resting state networks in the human brain. *Proc. Natl. Acad. Sci. U. S. A.* 104 (32), 13170.
- McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J., 1998. Analysis of fMRI data by blind separation into independent spatial components. *Hum. Brain Mapp.* 6 (3), 160–188.
- Mei, L., Figl, M., Rueckert, D., Darzi, A., Edwards, P., 2008. Statistical shape modelling: how many modes should be retained? *CVPRW* 1–8.
- Minka, T.P., 2001. Automatic choice of dimensionality for PCA. *Adv. Neural Inf. Process. Syst.* 598–604.
- Mohammadi, B., Kollwe, K., Samii, A., Krampfl, K., Dengler, R., Münte, T.F., 2009. Changes of resting state brain networks in amyotrophic lateral sclerosis. *Exp. Neurol.* 217, 147–153.
- Perlbarg, V., Marrelec, G., Doyon, J., Pelegrini-Issac, M., Lehericy, S., Benali, H., 2008. NEDICA: detection of group functional networks in fMRI using spatial independent component analysis. In *Proc. ISBI* 1247–1250.
- Perlbarg, V., Bellec, P., Anton, J.L., Pigrini-Issac, M., Doyon, J., Benali, H., 2007. Corsica: correction of structured noise in fMRI by automatic identification of ICA components. *Magn. Reson. Imaging* 25 (1), 35–46 Jan.
- Pinel, P., Thirion, B., Mriaux, S., Jobert, A., Serres, J., Le Bihan, D., Poline, J.B., Dehaene, S., 2007. Fast reproducible identification and large-scale databasing of individual functional cognitive networks. *BMC Neuroscience* 8 (1), 91 Oct.
- Rombouts, S.A.R.B., Damoiseaux, J.S., Goekoop, R., Barkhof, F., Scheltens, P., Smith, S.M., Beckmann, C.F., 2009. Model-free group analysis shows altered BOLD fMRI networks in dementia. *Hum. Brain Mapp.* 30 (1).
- Sadaghiani, S., Hesselmann, G., Kleinschmidt, A., 2009. Distributed and antagonistic contributions of ongoing activity fluctuations to auditory stimulus detection. *J. Neurosci.* 29 (42), 13410.
- Schwartzman, A., Dougherty, Y., Lee, Y., Gahremani, D., Taylor, J.E., 2009. Empirical null and false discovery rate analysis in neuroimaging. *Neuroimage* 44 (1), 71–82. doi:10.1016/j.neuroimage.2008.04.182 Jan.
- Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., Greicius, M.D., 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. *J. Neurosci.* 27 (9), 2349.
- Seeley, W.W., Crawford, R.K., Zhou, J., Miller, B.L., Greicius, M.D., 2009. Neurodegenerative diseases target large-scale human brain networks. *Neuron* 62 (1), 42–52.
- Shehzad, Z., Kelly, A.M., Reiss, P.T., Gee, D.G., Gotimer, K., Uddin, L.Q., Lee, S.H., Margulies, D.S., Roy, A.K., Biswal, B.B., et al., 2009. The resting brain: unconstrained yet reliable. *Cereb. Cortex*.
- Shmuel, A., Leopold, D.A., 2009. Neuronal correlates of spontaneous fluctuations in fMRI signals in monkey visual cortex: implications for functional connectivity at rest. *Hum. Brain Mapp.* 19, 2209–2229.
- Skudlarski, P., Jagannathan, K., Calhoun, V.D., Hampson, M., Skudlarska, B.A., Pearlson, G., 2008. Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations. *Neuroimage* 43 (3), 554–561.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., et al., 2004. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23, 208–219.
- Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F., 2009. Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U. S. A.* 106 (31), 13040–13045 Aug.
- Sorg, C., Riedel, V., Mühlau, M., Calhoun, V.D., Eichele, T., Laer, L., Drzezga, A., Forstl, H., Kurz, A., Zimmer, C., et al., 2007. Selective changes of resting-state networks in individuals at risk for Alzheimer's disease. *Proc. Natl. Acad. Sci. U. S. A.* 104 (47), 18760.
- Thirion, B., Dodel, S., Poline, J.B., 2006. Detection of signal synchronizations in resting-state fMRI datasets. *Neuroimage* 29 (1), 321–327 January.
- van den Heuvel, M.P., Mandl, R.C.W., Kahn, R.S., Hulshoff Pol, H.E., 2009. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Hum. Brain Mapp.* 30, 3127–3141.
- Wang, L., Zang, Y., He, Y., Liang, M., Zhang, X., Tian, L., Wu, T., Jiang, T., Li, K., 2006. Changes in hippocampal connectivity in the early stages of Alzheimer's disease: evidence from resting state fMRI. *Neuroimage* 31 (2), 496–504.