

Investigations into Resting-state Connectivity using Independent Component Analysis

FMRIB Technical Report TR05CB1
(A related paper has been accepted for publication in
Philosophical Transactions of the Royal Society,
Special Issue on 'Multimodal neuroimaging of brain connectivity')

Christian F. Beckmann, Marilena DeLuca, Joseph T. Devlin and Stephen M. Smith

Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB),
Department of Clinical Neurology, University of Oxford, John Radcliffe Hospital,
Headley Way, Headington, Oxford, UK
Corresponding author is Christian F. Beckmann: beckmann@fmrib.ox.ac.uk

Abstract

Inferring resting-state connectivity patterns from functional magnetic resonance imaging (FMRI) data is a challenging task for any analytical technique. In this paper we review a probabilistic independent component analysis (PICA) approach, optimised for the analysis of FMRI data (Beckmann and Smith, 2004), and discuss the role which this exploratory technique can take in scientific investigations into the structure of these effects. We apply PICA to FMRI data acquired at rest in order to characterise the spatiotemporal structure of such data, and demonstrate that this is an effective and robust tool for the identification of low-frequency resting-state patterns from data acquired at various different spatial and temporal resolutions. We show that these networks exhibit high spatial consistency across subjects and closely resemble discrete cortical functional networks such as visual cortical areas or sensory motor cortex.

Keywords: Functional Magnetic Resonance Imaging; brain connectivity; Resting-state fluctuations; Independent Component Analysis;

1 Introduction

Functional Magnetic Resonance Imaging (FMRI) has become an important neuroscientific tool for probing neural mechanisms in the human brain. Typical FMRI experiments have focused on the acquisition of T2*-sensitive MR images during periods of increased oxygen consumption (due to neuronal response to externally controlled experimental conditions) and contrast the measured image intensities with recordings obtained at 'rest'. Critically, some important quantitative concepts in FMRI analysis such as the calculation of percent signal change or the interpretation of de-activation implicitly hinge on a suitable definition of this baseline/rest signal. The baseline 'resting-state' of the brain itself, however, is a somewhat ill defined and poorly understood concept.

Of particular interest in this context are certain low-frequency fluctuations of the measured cerebral haemodynamics (around 0.01–0.1Hz) which exhibit complex spatial structure reminiscent of FMRI 'activation maps' and which can be identified in FMRI data taken both under rest condition and under external stimulation. Recently, some attention has been focused on the characterisation of these maps and the identification of possible origins of slow variations in the measured blood oxygen level dependent signal. Various researchers have suggested that these signal variations, temporally correlated across the brain, are of neuronal origin and correspond to functional *resting-state networks* (RSNs) which jointly characterise the neuronal baseline activity of the human brain in the absence of deliberate and/or externally stimulated neuronal activity, and may reflect functionally distinct networks.

Biswal et al. (1995) first demonstrated the feasibility of using FMRI to detect such spatially distributed networks within primary motor cortex during resting-state by calculating temporal correlations across the brain with the time course from a seed voxel whose spatial location was chosen from a prior finger-tapping study. The temporal signal from a seed voxel in the motor cortex was correlated with other motor cortex voxels and uncorrelated with other voxels, with major frequency peaks in

the resting correlations at around 0.02Hz. Lowe et al. (1998) found similar results using both single-slice low time of repetition (TR of 130ms) and whole-head volumes with longer TR (2000ms) while Xiong et al. (1999) describe functional connectivity maps that cover additional non-motor areas. Based also on findings from PET studies, the existence of a *default mode brain network* involving several regions including the posterior cingulate cortex has been proposed (Raichle et al., 2001; Shulman et al., 1997; Mazoyer et al., 2001). Using simultaneously acquired EEG and FMRI data under rest, Goldman et al. (2002) have shown that the variation in Alpha rhythm in EEG (8-12Hz) is correlated with the FMRI measurements. In particular, the authors report that increased alpha power was correlated with decreased BOLD signal in multiple regions of occipital, superior temporal, inferior frontal, and cingulate cortex, and with increased signal in the thalamus and insula. These results have important implications for interpretation of RSNs as they suggest a neuronal cause for these fluctuations.

Alternatively, it has been argued that these effects simply reflect vascular processes unrelated to neuronal function, which would make RSNs of less interest to neuroscience (though still of potential clinical interest). Physiological noise in the resting brain and its echo-time and field strength dependencies were investigated by Kruger and Glover (2001) who showed that physiological noise demonstrates a field strength dependency, exceeds the thermal as well as scanner noise at 3T and is increased in grey matter (see also Woolrich et al. (2001)). Various researchers have investigated the relation between low-frequency fluctuations in the measured BOLD signal and other physiological observations: Obrig et al. (2000) reviewed and studied low-frequency variations in oxygenation, cerebral blood flow (CBF) and metabolism and report significant correlations with similar fluctuations observed by near infrared spectroscopy (NIRS). More recently, Wise et al. (2004) have investigated the influence of arterial carbon dioxide fluctuations by using the endtidal level of exhaled carbon dioxide as covariate of interest in a General Linear Model (GLM) analysis. The most significant changes were concentrated in the occipital, parietal and temporal lobes as well as in the cingulate cortex, and suggest that vascular processes (unrelated to neuronal function) play a significant role in the generation of such resting-state patterns.

Estimating the temporal and spatial characteristics of these low-frequency fluctuations from FMRI data presents a formidable challenge to analytical techniques. In the majority of existing studies, resting patterns are inferred by a correlation analysis of the voxel-wise FMRI recordings against a reference time course obtained from secondary recordings (e.g. from EEG, NIRS or physiologic measurements like the carbon-dioxide concentration) or simply by regressing against a single voxel's time course from resting data which is believed to be of functional relevance (seed-voxel based correlation analysis). These techniques fundamentally test very specific hypotheses about the temporal structure of these effects. Recently, however, Independent Component Analysis has successfully been applied to the estimation of certain low-frequency patterns (Goldman and Cohen, 2003; Kiviniemi et al., 2003; Greicius et al., 2004). An important benefit of such exploratory techniques over more hypothesis-based techniques is the ability to identify various types of signal fluctuations by virtue of their spatial and/or temporal characteristics without the need to specify an explicit temporal model. Such flexibility in data modelling is essential in cases where the effects of interest are not very well understood and cannot be predicted accurately.

This paper is organised as follows: in section 2 we review a probabilistic approach to Independent Component Analysis (PICA) specifically optimised for the analysis of FMRI data (Beckmann and Smith, 2004). Section 3 discusses the constraints of this exploratory data analysis technique when used for the identification of large-scale noise fluctuations. In particular, we demonstrate that optimisation for maximally independent spatial sources does not imply an inability to estimate largely overlapping spatial maps. We demonstrate the ability of PICA to extract resting fluctuations and apply the technique to FMRI resting data in order to test a set of important hypotheses about the structure of resting-state connectivity in the human brain. In particular, we will investigate (i) if and how estimated source processes are driven by less interesting physiological effects such as the cardiac or respiratory cycle, (ii) the spatial characteristics of estimated maps in terms of locality within grey matter and (iii) the consistency of maps obtained from multiple subjects.

2 Decomposing FMRI data using ICA

Independent Component Analysis (ICA, Comon (1994); Bell and Sejnowski (1995); McKeown et al. (1998)) is a technique which decomposes a 2-dimensional (time \times voxels) data matrix¹ into a set of time courses and associated spatial maps which jointly describe the temporal and spatial characteristics of underlying hidden signals (components). A probabilistic ICA model extends this by assuming that the p -dimensional vectors of observations (time series in the case of FMRI data) is generated from a set of $q (< p)$ statistically independent non-Gaussian sources (spatial maps) via a linear and instantaneous 'mixing' process corrupted by additive Gaussian noise $\eta(t)$:

$$\mathbf{x}_i = \mathbf{A}\mathbf{s}_i + \boldsymbol{\eta}_i \tag{1}$$

¹Here, we only discuss the case of a decomposition into spatially independent source signals; the reason for this will become apparent later.

Here, x_i denotes the individual measurements² at voxel location i , s_i denotes the non-Gaussian source signals contained in the data and η_i denotes Gaussian noise³ $\eta_i \sim \mathcal{G}(0, \sigma^2 \Sigma_i)$.

The $p \times q$ dimensional mixing matrix \mathbf{A} is assumed to be non-degenerate, i.e. of rank q . Solving the blind separation problem requires finding a linear 'unmixing' matrix \mathbf{W} of dimension $q \times p$ such that

$$\hat{\mathbf{s}} = \mathbf{W}\mathbf{x}$$

is a good approximation to the true source signals \mathbf{s} .

The PICA model is similar to the standard GLM with the difference that, unlike the design matrix in the GLM, the mixing matrix \mathbf{A} is no longer pre-specified prior to model fitting but will be estimated from the data. The spatial source signals correspond to parameter estimate images in the GLM with the additional constraint of being statistically independent of each other.

2.1 Parameter estimation

Without loss of generality we can assume that the source signals have unit variance. If the noise covariance Σ_i is known, we can pre-whiten the data and obtain a new representation $\bar{x}_i = \bar{\mathbf{A}}s_i + \bar{\eta}_i$, where $\bar{\eta}_i \sim \mathcal{G}(0, \sigma^2 \mathbf{I})$, i.e. where the noise covariance is isotropic at every voxel location. To simplify notation, we will henceforth assume isotropic noise and drop the additional bar.

Noise and signal are uncorrelated, so the data covariance matrix $\mathbf{R}_x = \langle x_i x_i^t \rangle = \mathbf{A}\mathbf{A}^t + \sigma^2 \mathbf{I}$, i.e. the unknown mixing matrix \mathbf{A} can be estimated as the matrix square root of $\mathbf{R}_x - \sigma^2 \mathbf{I}$: let \mathbf{X} be a $p \times N$ matrix containing all N different fMRI time series in its columns and let $\mathbf{X} = \mathbf{U}(N\mathbf{\Lambda})^{\frac{1}{2}}\mathbf{V}$ be the singular value decomposition of \mathbf{X} . Then

$$\hat{\mathbf{A}}_{\text{ML}} = \mathbf{U}_q(\mathbf{\Lambda}_q - \sigma^2 \mathbf{I}_q)^{\frac{1}{2}}\mathbf{Q}^t, \quad (2)$$

where \mathbf{U}_q and $\mathbf{\Lambda}_q$ contain the first q Eigenvectors and Eigenvalues. The matrix \mathbf{Q} denotes a $q \times q$ orthogonal rotation matrix, i.e. a matrix with $\mathbf{Q}\mathbf{Q}^t = \mathbf{I}$. This matrix is not directly identifiable from the data covariance matrix since \mathbf{R}_x is invariant under post-multiplication of \mathbf{A} by any orthogonal rotation $\bar{\mathbf{Q}}$ given that $(\mathbf{A}\bar{\mathbf{Q}})(\mathbf{A}\bar{\mathbf{Q}})^t = \mathbf{A}\bar{\mathbf{Q}}\bar{\mathbf{Q}}^t\mathbf{A}^t = \mathbf{A}\mathbf{A}^t = \mathbf{R}_x - \sigma^2 \mathbf{I}$.

Estimating the mixing matrix \mathbf{A} , however, reduces to identifying the square matrix \mathbf{Q} after whitening the data with respect to the noise covariance Σ_i and projecting the temporally whitened observations onto the space spanned by the q Eigenvectors of \mathbf{R}_x with largest Eigenvalues. The maximum likelihood estimates of sources and σ are obtained using generalised least squares:

$$\begin{aligned} \hat{\mathbf{s}}_{\text{ML}} &= \hat{\mathbf{W}}\mathbf{x} \quad \text{with} \quad \hat{\mathbf{W}} = (\hat{\mathbf{A}}^t \hat{\mathbf{A}})^{-1} \hat{\mathbf{A}}^t \\ \text{and} \quad \hat{\sigma}_{\text{ML}}^2 &= \frac{1}{p-q} \sum_{l=q+1}^p \lambda_l. \end{aligned} \quad (3)$$

Solving the model in the case of an *unknown* noise covariance can be achieved by iterating estimates of the mixing matrix and the sources and re-estimating the noise covariances from the residuals $\hat{\eta}$. The form of Σ_i typically is constrained by a suitable parameterisation; here we use the common approaches to fMRI noise modelling (Bullmore et al., 1996; Woolrich et al., 2001), and restrict the structure to autoregressive noise. However, since the exploratory approach allows modelling of various sources of variability, e.g. temporally consistent physiological noise, as part of the signal in equation 1, the noise model itself can actually be quite simplistic.

A consequence of the isotropic noise model is that as an initial pre-processing step we will modify the original data time courses to be normalised to zero mean and unit variance. This pre-conditions the data under the null hypothesis of no signal: the data matrix \mathbf{X} is identical (up to second order statistics) to a simple set of realisations from a $\mathcal{G}(0, \mathbf{I})$ noise process. Any signal will have to reveal itself via its deviation from Gaussianity.

The maximum likelihood estimators depend on knowledge of the number of underlying sources q . In the noise free case this quantity can easily be deduced from the rank of the covariance of the observations \mathbf{R}_x which is of rank q . In the presence of isotropic noise, however, the covariance matrix will be of full rank where the additional noise has the effect of raising the Eigenvalues of the covariance matrix by σ^2 (Roberts and Everson, 2001). Inferring the number of estimable source processes amounts to testing for sphericity of Eigenspaces beyond a given threshold level (Beckmann and Smith, 2004). Simplistic criteria like the reconstruction error or predictive likelihood will naturally predict that the accuracy steadily increases with increased dimensionality. Thus, criteria like retaining 99.9% of the variability result in arbitrary threshold levels (Beckmann

²For simplicity we assume de-meaned data.

³The covariance of the noise is allowed to be voxel dependent in order to encode the vastly different noise covariance observed within different tissue types (Woolrich et al., 2001).

et al., 2001). This problem is intensified by the fact that the data covariance \mathbf{R}_x is being estimated by the sample covariance matrix. In the absence of any source signals, the Eigenspectrum of this sample covariance matrix is not identical to $\sigma^2 \mathbf{I}_p$ but instead distributed skewed around the true noise covariance: the Eigenspectrum will depict an apparent difference in the significance of individual directions within the noise (Everson and Roberts, 2000), even in the absence of signal. In the case of Gaussian noise, however, this 'skew' of the Eigenspectrum is of analytic form: the Eigenvalues have a Wishart distribution and we can adjust the observed Eigenspectrum by the quantiles of the predicted cumulative distribution of Eigenvalues from Gaussian noise (Johnstone, 2000), prior to estimating the model order. If we assume that the source distributions $p(s)$ are Gaussian, the model then reduces to probabilistic PCA (Tipping and Bishop, 1999) and we can use Bayesian model selection criteria. Within the PICA approach, we use the Laplace approximation to the posterior distribution of the model evidence that can be calculated efficiently from the adjusted Eigenspectrum (Minka, 2000; Beckmann and Smith, 2004).

In order to complete the estimation of the mixing matrix and the sources, we need to optimise an orthogonal rotation matrix \mathbf{Q} in the space of whitened observations:

$$\hat{\mathbf{s}} = \mathbf{W}\mathbf{x} = \mathbf{Q}\tilde{\mathbf{x}}, \quad (4)$$

where $\tilde{\mathbf{x}} = (\mathbf{\Lambda}_q - \sigma^2 \mathbf{I}_q)^{-1/2} \mathbf{U}_q^t \mathbf{x}$ denotes the spatially whitened data.

Hyvärinen and Oja (1997) have presented an elegant fixed point algorithm that uses approximations to neg-entropy in order to optimise for non-Gaussian source distributions and give a clear account of the relation between this approach to statistical independence. In brief, the individual sources are obtained by projecting the data \mathbf{x} onto the individual rows of \mathbf{Q} , i.e. the r th source is estimated as

$$\hat{\mathbf{s}}_r = \mathbf{v}_r^t \tilde{\mathbf{x}},$$

where \mathbf{v}_r^t denotes the r th row of \mathbf{Q} . In order to optimise for non-Gaussian source estimates, Hyvärinen and Oja (1997) propose the following contrast function:

$$J(\mathbf{s}_r) \propto [\mathcal{E}\{F(\hat{\mathbf{s}}_r)\} - \mathcal{E}\{F(\nu)\}], \quad (5)$$

where ν denotes a standardised Gaussian variable, \mathcal{E} denotes the expectation and F is a general non-quadratic function that combines the high-order moments of \mathbf{s}_r in order to estimate the amount of non-Gaussianity in the individual sources. From equation 5, the vectors \mathbf{v}_r^t are optimised to maximise $J(\hat{\mathbf{s}}_r)$ using an approximative Newton method (Hyvärinen and Oja, 1997).

2.2 Inference

After estimating the mixing matrix $\hat{\mathbf{A}}$, the source estimates are calculated by projecting each voxel's time course onto the time courses contained in the columns of the unmixing matrix $\hat{\mathbf{W}}$. In the case where the model order q was estimated correctly, the estimated noise is a linear projection of the true noise and is unconfounded by residual signal. At every voxel location we have pre-conditioned the data such that \mathbf{x}_i has unit standard deviation and the estimate of the noise variance $\hat{\sigma}_i^2$ at each voxel location will approximately equal the true variance of the noise. We can thus convert the individual spatial IC maps \mathbf{s}_r . into 'Z-statistic maps' \mathbf{z}_r . by dividing the raw IC maps by the standard error of the residual noise.

In order to assess the Z-maps for 'significantly activated' voxels, we employ mixture modelling of the probability density of the Z-statistic spatial maps.

From equation 3 it follows that $\hat{\mathbf{s}}_i = \hat{\mathbf{W}} \mathbf{A} \mathbf{s}_i + \hat{\mathbf{W}} \boldsymbol{\eta}_i$, i.e. the noise term in equation 1 manifests itself as additive Gaussian noise in the estimated sources. We therefore model the distribution of the spatial intensity values of each Z-map by a mixture of one Gaussian and two Gamma distributions, to model background noise and positive and negative BOLD effects (Hartvig and Jensen, 2000; Beckmann et al., 2003). The mixture is fitted using an expectation-maximisation algorithm Dempster et al. (1977). In cases where the number of 'active' voxels is very small, the relative proportions of the Gamma densities in the overall mixture distribution might be estimated as zero. In this case, a simple transformation to spatial Z-scores and subsequent thresholding is appropriate, i.e. reverting to null-hypothesis testing instead of the otherwise preferable alternative-hypothesis testing. Otherwise we can evaluate the fitted mixture model to calculate the posterior probability of 'activation' as the ratio of the probability of intensity value under the 'noise' Gaussian relative to the sum of probabilities of the value under the 'activation' Gamma densities⁴.

Any threshold level, though arbitrary, directly relates to the loss function we like to associate with the estimation process, e.g. a threshold level of 0.5 places an equal loss on false positives and false negatives (Hartvig and Jensen, 2000).

2.3 PICA Algorithm Overview

The individual steps that constitute the Probabilistic Independent Component Analysis are illustrated in figure 1. The demeaned original data is first normalised to unit variance at each voxel location. If appropriate spatial information is available,

⁴In this case 'activation' is to be understood as signal that 'cannot be explained as random correlation coefficient'.

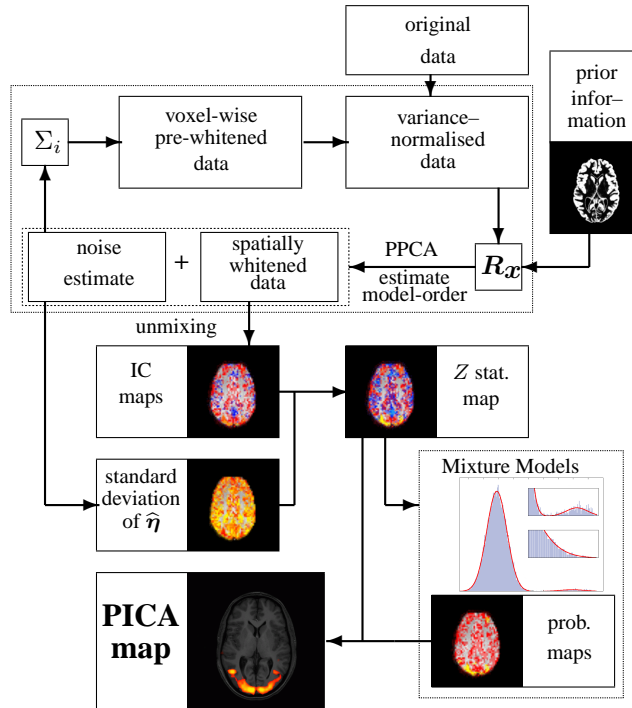


Figure 1: Schematic illustration of the probabilistic ICA model (Beckmann and Smith, 2004)

this is encoded in the estimation of the sample covariance matrix R_x . Individual voxel weights, e.g. gray-matter segmentation, can be used to calculate a weighted covariance matrix while voxel-pair weightings can be used to calculate the within-group covariance (Beckmann and Smith, 2004). Probabilistic PCA is used to infer upon the unknown number of sources and results in an estimate of the noise and a set of spatially whitened observations. We can estimate the noise covariance structure Σ_i from the residuals in order to voxel-wise (temporally) pre-whiten and re-normalise the data and iterate the entire cycle. Estimation of Σ_i from residuals in the case of autocorrelated noise can be achieved as described by Woolrich et al. (2001). In practice, the output results do not suggest a strong dependency on the form of Σ_i and preliminary results suggest that it is sufficient to iterate these steps only once. From the spatially whitened observations, the individual component maps are obtained using a modified fixed point iteration scheme (FastICA (Hyvärinen and Oja, 1997)) to optimise for non-Gaussian source estimates via maximising the neg-entropy. These maps are separately transformed to Z -scores. In contrast to raw IC estimates which only encode the estimated signal, these Z score maps depend on the amount of variability explained by the entire decomposition at each voxel location relative to the residual noise similar to statistical parametric maps from a GLM analysis. This is an important aspect of the probabilistic ICA model as now these maps also reflect the degree to which the signal explained within this model fits to the data and, unlike standard ICA, no longer ignores the signal variation which remains unaccounted for. Finally, Gaussian/Gamma Mixture Models are fitted to the individual Z maps in order to infer voxel locations that are significantly modulated by the associated time course.

3 Estimating overlapping maps using ICA

The choice of optimising for independence between spatial maps could equally well be replaced by optimising for independence between time courses. Different authors have argued in favour of one or the other technique, where the main objection appears to revolve around the question of whether orthogonality (i.e. uncorrelatedness) between estimated sources should be enforced in the temporal or spatial domain (Friston, 1998; Petersen et al., 2000). At a conceptual level, the notion of orthogonality is overly restrictive in either domain: for temporal modes, the existence of stimulus correlated effects (e.g. motion artefacts or higher order brain function) means that enforced orthogonality necessarily results in a mis-representation of underlying temporal signals. Similarly, for spatial modes, (Friston, 1998) has argued that even though different brain function might be spatially localised, the principle of 'functional integration' might imply that neuronal processes share a large proportion of cortical anatomy. These arguments suggest that independence and implied orthogonality are always suboptimal for the analysis of data which is as complicated as that obtained from functional MRI experiments.

From a signal detection point of view, however, it is important to consider the extent to which signal 'appears' in space or




Figure 2: Estimation of largely overlapping signals in the presence of noise: 2 source signals with spatial correlation of $\rho \sim 0.5$ were introduced into Gaussian noise ($\sigma = 3$) to form an artificial data set of size 250×10000 (a). In the presence of noise, the spatial correlation of least-squares estimates is significantly reduced (b). Estimating principal components from the data results in a poor representation in the temporal and spatial domain (c). ICA estimates from the same data (d) show much improved detection power and represent the spatial maps and time courses well. Thresholded maps (e) are again highly correlated at $\rho \sim 0.47$. Note that spatial ICA maps, like PCA maps, are constrained to be orthogonal. This restriction itself, therefore, does not necessarily imply poor spatial representation of signals even in cases where 'true' spatial maps are highly correlated.

time. Within the temporal domain, signal often spans the entire length of an experiment. If the 'true' temporal characteristics of different signals are partially correlated (e.g. stimulus-correlated motion), a decomposition which enforces orthogonality in the temporal domain will necessarily mis-represent at least one of the time series in order to satisfy the constraint. In the spatial domain, however, 'signals' in fMRI are sparse and typically are contained in a small proportion of all voxels. Even for what in fMRI are considered 'large' activation clusters or for artefactual sources with large spatial extent (e.g. image ghosts), only a fraction of intra-cranial voxels are involved⁵. In the presence of noise, the majority of voxels in any spatial maps have random 'background noise' value and will reduce the observed spatial correlation such that even when 'true' spatial maps are significantly overlapping, a decomposition which enforces orthogonality between estimated spatial maps can still give a relatively accurate representation of the signal.

Formally, consider the case of two source signals s_1 and s_2 , represented as column vectors of length N , and (zero-mean) Gaussian noise η_1 and η_2 with variance σ_1^2 and σ_2^2 . In the presence of noise, the correlation changes from

$$\rho(s_1, s_2) = \frac{s_1^t s_2}{N \sqrt{\text{Var}(s_1)} \sqrt{\text{Var}(s_2)}}$$

⁵For residual motion artefacts, every voxel is theoretically influenced by an associated motion time series, but only voxels near intensity boundaries are detectable.

to

$$\rho(s_1 + \eta_1, s_2 + \eta_2) = \frac{s_1^t s_2}{N \sqrt{\text{Var}(s_1) + \sigma_1^2} \sqrt{\text{Var}(s_2) + \sigma_2^2}}.$$

When signals are sparse, $\text{Var}(s_1)$ and $\text{Var}(s_2)$ are small and the denominator of $\rho(s_1 + \eta_1, s_2 + \eta_2)$ is dominated by the noise variance. The reduction in correlation between the noise-free and noisy case is a function of the signal amplitude modulation, the sparseness and the relative noise level.

As an example figure 2(a) shows two partially overlapping spatial 'signals' each occupying $\sim 17\%$ of the total image areas together with two artificial time courses. Due to their partial overlap, these source signals are spatially correlated with a correlation of $\rho \sim 0.5$. In the absence of noise these maps can not be estimated accurately by any technique enforcing orthogonality between estimated spatial maps. In the presence of noise⁶, however, the spatial correlation between linear estimates reduces significantly: figure 2(b) shows the spatial maps obtained from performing linear regression of the data against the 'true' time series. The spatial maps obtained from a PCA decomposition (c) have ~ 0 spatial correlation and fail to identify the 'true' spatial maps. Also, the temporal characteristics of the signal are not well represented. By comparison, the estimated spatial maps from an ICA decomposition (d) well represent signal in space and time. Although the spatial sources are clearly visible, the spatial correlation between the estimated spatial maps is still ~ 0 . This is a consequence of the optimisation for maximally non-Gaussian source projections. Final thresholded ICA maps derived from a Gaussian/Gamma mixture model on the noisy maps give a reasonably good spatial representation for the original sources: the estimated thresholded maps (e) have large spatial correlation ($\rho \sim 0.47$).

This example demonstrates that the mathematical constraint of orthogonality within the set of spatial maps does not necessarily imply that large areas of 'activation' which overlap significantly between maps can no longer be extracted. Instead, the amount to which this mathematical constraint restricts the estimation of partly overlapping sources is a function of (i) the overall sparseness of signals and (ii) the signal-to-noise ratio. This suggests that in practice the constraints induced by optimising for independence are less restrictive in the spatial domain than the temporal domain. Though compensating for partial correlation of 'signal' by anti-correlating 'noise' conceptually is also possible in the temporal domain, the significantly lower number of time points does not typically provide a sufficient number of 'background' time points that could be utilised to ensure orthogonality while not altering 'interesting' portions of the estimated source signals. This property is particular important for investigating resting-state networks because it means that functionally distinct systems can overlap anatomically as long as they have sufficiently distinct time courses.

4 Experimental Method

In order to characterise the low-frequency structured noise components in 'resting' data, we acquired different data sets to address four specific questions:

1. To compare seed-voxel correlation techniques with PICA, we collected 200 volumes from a single subject under rest and active finger tapping (30s on/off block design). Data were acquired on a Phillips NT 1.5T MRI system with a notional 2x2x8mm resolution, a repetition time (TR) of 3s and an echo-time (TE) of 40ms.
2. To evaluate the extent to which neural effects can be distinguished from non-neural physiological effects such as aliased cardiac or respiratory cycles, we collected resting data with a high temporal resolution (B0=3T, TR=125ms, TE=30ms, 3.75x3.75mm in-plane resolution). The data consists of 2160 slices through a single axial plane chosen to intersect the sensori-motor cortices bilaterally. In addition, we collected 60 volumes under a 30s on/off bilateral finger tapping paradigm at a typical TR of 3s. All data were acquired on a Varian-Siemens 3T MRI system.
3. To determine whether low frequency resting fluctuations appear within grey matter or are instead driven by contributions from larger blood vessels, we collected 300 volumes (12 slices) of resting data at 3T with spatial resolution of 2x2x6mm and TR=3s.
4. Finally, to investigate the spatial consistency of resting-state patterns across subjects, data were collected from nm 10 subjects during rest. For each, 200 volumes of whole head functional data were acquired at 3T with typical fMRI resolution (3x3x3mm, TR=3.4s, TE=40ms). In addition, a high-resolution T1-weighted reference scan (1x1x1.5mm resolution) was also acquired for the purpose of anatomical localisation.

Subjects were lying supine in the MRI scanner and instructed to keep their eyes closed and not to fall asleep during functional scanning.

⁶Zero-mean Gaussian noise with $\sigma = 3$; the maximum signal amplitude modulation is 2.

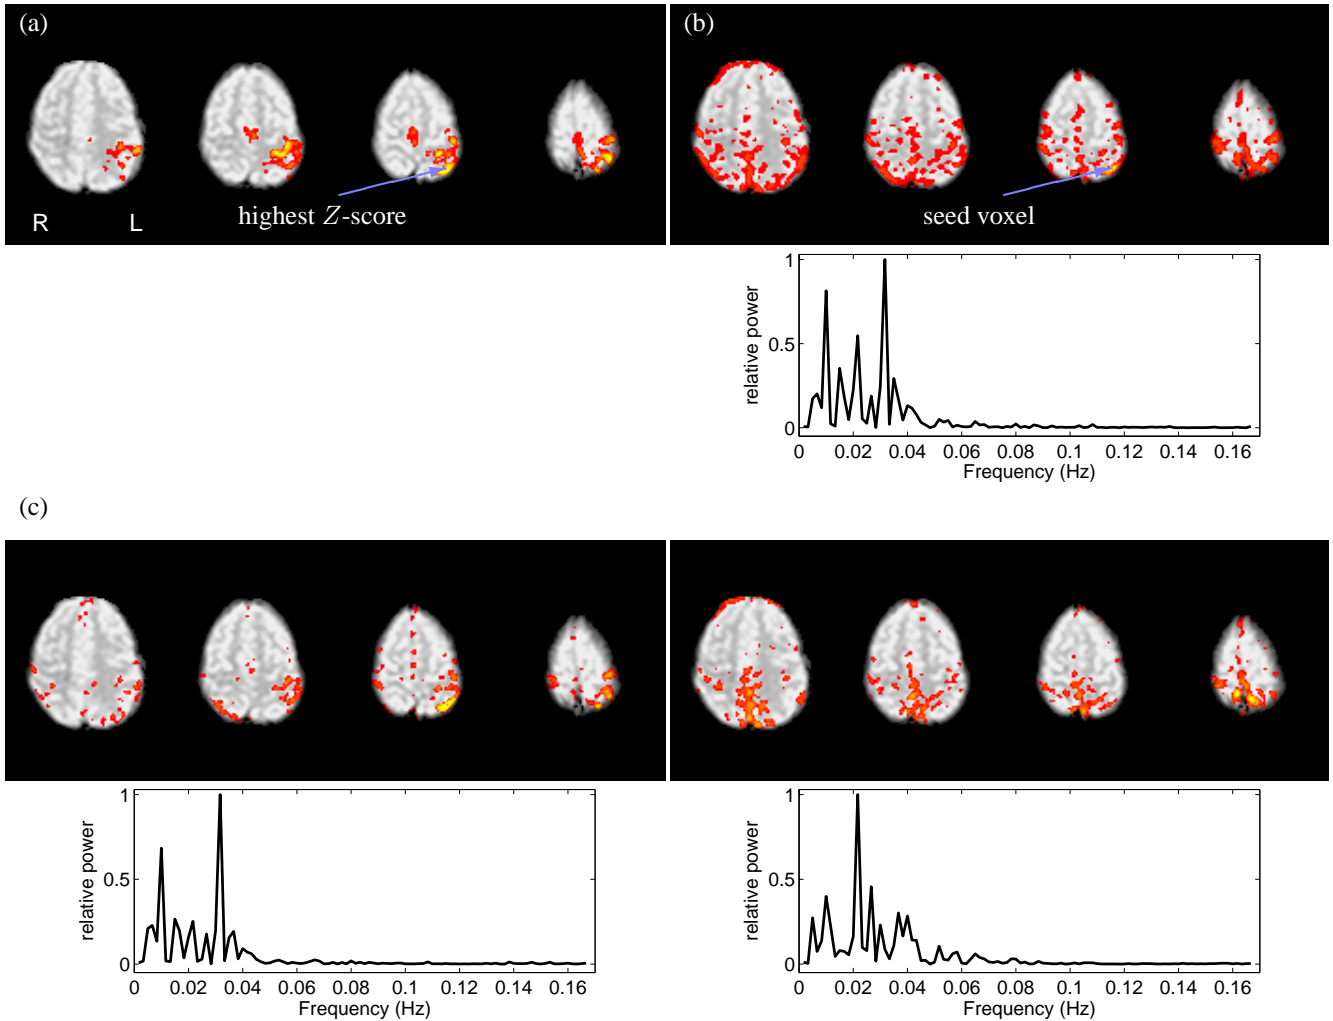


Figure 3: Comparison of seed-voxel based correlation analysis and PICA: for seed-voxel based analysis, an activation data set is first analysed to identify the location of most significant response to external stimulation (a). The time course of the coinciding voxel in the resting data is used as a reference time course for a correlation analysis of all other time courses acquired under rest. The resulting correlation map (b) shows significant resting correlation in similar motor areas as the activation map but also identifies part of the ipsi-lateral motor cortex. In addition, other regions outside of these motor areas were also found including medial and lateral posterior parietal areas and prefrontal regions. In a PICA decomposition of the resting data (c), similar cortical regions are identified by two (out of 40) separate spatial maps. The multiple regression framework implicit in a PICA decomposition separates resting correlations in motor areas (left) from other cortical areas. This separation is induced by the fact that the associated time courses are significantly different: the associated normalised power spectra show different peak frequencies. All spatial maps were thresholded using mixture modelling (at $p > 0.5$) and are shown in radiological convention.

4.1 Analysis Methods

The individual data sets were preprocessed before running correlation-based or ICA-based statistical analyses using tools from the FMRIB Software Library (FSL, www.fmrib.ox.ac.uk/fsl). Time series were first realigned to correct for small head movements (Jenkinson et al., 2002). Then non-brain (e.g. scalp and CSF) were removed using an automated brain extraction tool (Smith, 2002). Finally, the data were spatially smoothed using a Gaussian kernel (5mm). After statistical analysis (whether correlation- or ICA-based), the resulting statistical maps were thresholded using histogram mixture modelling as described above.

4.2 Seed-voxel based correlation analysis vs. multivariate PICA regression

Activation-seeded correlation analysis is based on the hypothesis that in resting data the low-frequency temporal fluctuations are correlated in regions which are functionally associated (Biswal et al., 1995; Lowe et al., 1998). In this approach an activation dataset (e.g. under a simple motor paradigm) is acquired along with data at rest. The activation data is first analysed to identify areas which activate significantly. The coordinates of the highest activating voxel are then used to define a seed voxel in the resting data. A statistical map is generated by calculating the correlation of all time courses in the resting data against the time course of the seed voxel in order to find a temporally consistent resting network. The applicability of this technique, however, is limited by the fact that seed-voxel based analysis relies on the time course at the seed voxel location being a 'good' representative for the set of correlated voxels under rest. As a consequence, the seed-voxel based approach is restricted to cases where seed areas can be inferred accurately and robustly from activation studies (like motor area). Furthermore, the choice of the seed voxel is rather arbitrary (as, indeed, is the exact location of a peak Z-stat) and can be biased by different types of fMRI noise. In particular, the usefulness of such a correlation analysis is severely limited in cases where the reference time course itself is a mixture of time courses, e.g. different low-frequency fluctuations, spatially structured high-frequency signals such as that induced by the $N/2$ ghost, head motion etc. These problems are analogous to those that characterise the difference between simple correlation analysis and the GLM for model-based fMRI analysis: a multiple regression model can account for temporal effects which coincide at a single voxel location.

Figure 3 illustrates the difference between seed-voxel based correlation analysis and PICA using data from a simple finger tapping experiment and data acquired under rest (i.e. the first dataset). Model-based analysis of the activation data produced plausible motor cortex activation in the contra-lateral hemisphere (a), although in this case, the peak Z-score was located in the post-central gyrus rather than in motor cortex. Using this voxel as the seed time course in a subsequent correlation-based analysis of the resting data (i.e. excluding the motor task blocks) shows significant correlation in similar areas such as the motor cortex bilaterally and the supplemental motor area (SMA) along the midline (b). On the other hand medial posterior cortical areas and frontal parts also show significant correlation, despite not being identified as parts of the motor system engaged in the finger tapping contrast (a). Based on the Laplace approximation of the Bayesian model evidence, the PICA approach estimates 40 components, including various artefacts such as Nyquist ghosting, head-motion and large blood vessels. Two of the remaining components (shown in figure 3(c)) jointly cover almost identical post-thresholded areas as the map obtained from seed-voxel based correlation analysis. Within the PICA approach, these areas are separated into different spatial maps due to the fact that the associated time courses are sufficiently different. (as can be seen by the different power spectra). The PICA decomposition suggests that the voxels shown in figure 3(b) are part of two different spatial patterns which appear in the single correlation map by virtue of the fact that the seed voxel has partial correlation with voxels shown in the two PICA-derived maps. The multiple regression analysis underlying a PICA decomposition, by comparison, can separate these effects and gives a more plausible representation of a motor network.

4.3 Relation between physiological noise and resting-state fluctuations

One question that arises with respect to resting-state networks such as the one shown in figure 3 is whether the findings are functionally significant, or whether they are simply a consequence of aliased physiological effects such as the cardiac and respiratory cycles. In order to investigate the frequency characteristics of resting fluctuations we acquired single slice data covering the motor cortex at low TR (125ms) and at a more typical TR (3s). The high temporal sampling data is necessary to separate low-frequency effects from signal fluctuations due to the cardiac or respiratory effects. These occur naturally at frequencies of 0.3-1Hz and so can easily become aliased by normal TR sampling to give significant power at the low frequencies (0.015 to 0.035Hz) typical of resting-state fluctuations. By collecting low TR data such aliasing will be avoided. This data has about twice as many samples in the temporal domain than voxels across space. In order to reduce computational load, therefore, we assumed a block-diagonal form of the data covariance matrix for the initial PCA dimensionality reduction which is part of the spatial PICA decomposition.

Figure 4 shows PICA estimates obtained from low-TR resting data (left) and analysis results from activation data (self-paced bilateral finger tapping) acquired at a more typical sampling rate of 3s (right). On the low-TR data, the separate components found include a single cardiac-cycle-related map with peak frequency of ~ 1 Hz (fig 4(a)), a respiration-related map with peak frequency of ~ 0.3 Hz (fig. 4(b)) as well as a map showing the spatial extent of lower frequency resting fluctuations (0.02-0.1Hz, fig 4(c)) - in this case largely contained within sensori-motor areas. With high temporal sampling the PICA approach clearly separates simple physiological noise components from resting fluctuations. By comparison, figure 4(d-f) shows corresponding PICA maps estimated from data acquired with a more typical TR. Here, the spatial maps of the respiratory and cardiac fluctuations were identified purely based on their spatial correspondence with maps shown in figure 4(a,b) (spatial correlation of 0.64 and 0.42 respectively). At this more normal temporal sampling rate, the simple frequency characteristics of these effects are no longer detectable due to aliasing. The primary activation map, however, relates to activation due to a 30s on/off block design (periodicity of 0.01677) and can be identified easily both in the spatial and frequency domain. This

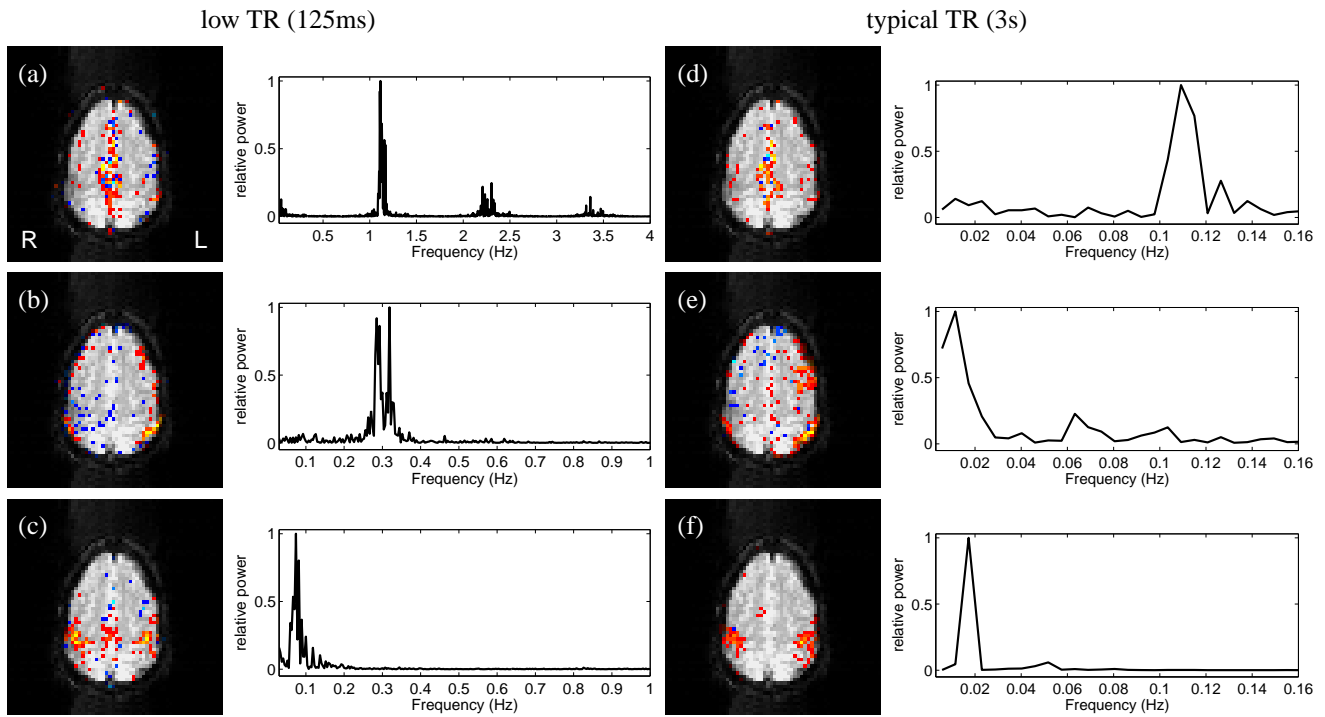


Figure 4: Investigating the relationship between physiological noise and low-frequency fluctuations in the spatial and temporal domain: The different components estimated from the low-TR data (left) show a clear separation of physiological artefacts induced by the cardiac cycle (a) and the respiratory cycle (b) from low-frequency fluctuations (c) both in the spatial maps and the corresponding power spectra. At higher TR the temporal signature of the cardiac and respiratory cycles become aliased and no longer identifiable in the frequency domain. The spatial maps (d,e), however, show a high degree of correspondence with maps (a) and (b) (spatial correlation of 0.64 and 0.42, respectively), suggesting that the PICA approach is able to separate relatively uninteresting physiological noise from other effects such as resting-state maps even in cases where the physiological noise fluctuations become aliased in the temporal domain.

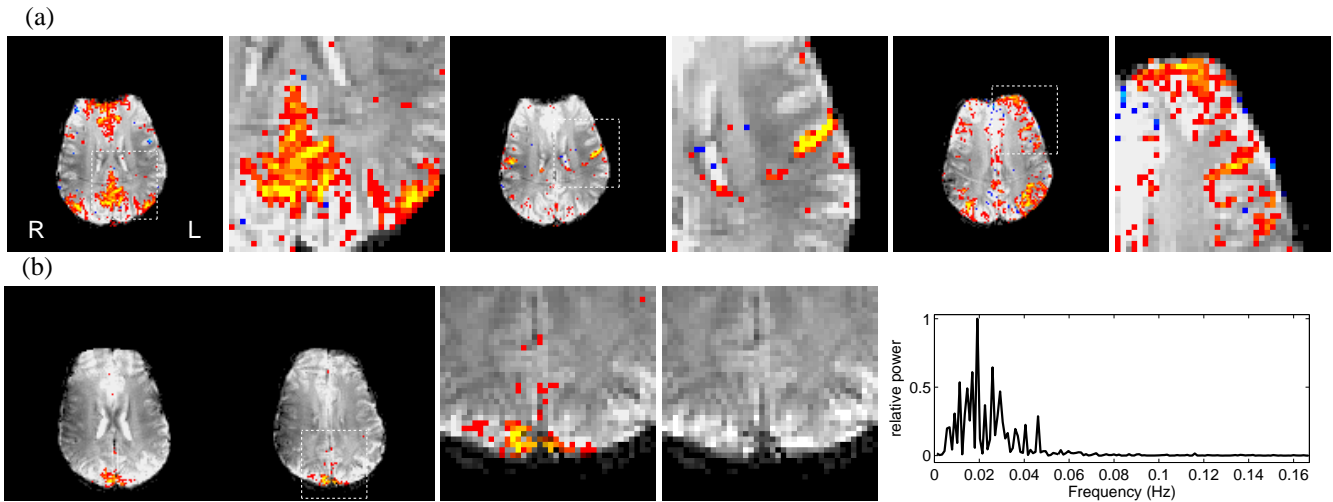


Figure 5: Investigating the spatial structure of resting-state fluctuation: PICA analysis of EPI data acquired at 2x2 mm resolution in the x-y plane suggests that the resting-state fluctuations are well localised in grey-matter (a). Furthermore, they appear to be spatially different from 'blood vessel networks' (BVNs) which mainly show larger blood vessels and surrounding tissue (b).

suggests that even at higher TR PICA-derived spatial maps separate cardiac and respiratory effects from other effects such as activation maps or low-frequency resting-state fluctuations.

4.4 Spatial characteristics of resting-state fluctuations

The functional relevance of low-frequency patterns depends on the spatial locality in grey-matter. At typical spatial fMRI resolution, the estimated low-frequency patterns indeed appear to be 'grey matter networks'. Figure 5(a) shows example spatial maps found using PICA on the data with higher in-plane spatial resolution (2x2mm; third data set). The significant voxels do generally lie within grey matter and include little or no white matter. In almost all cases a PICA decomposition also generates spatial maps depicting 'blood vessel networks' (BVNs). These have a similar power spectrum with peak frequencies of around 0.2Hz but mainly show larger blood vessels and surrounding tissue (fig. 5(b)).

4.5 Consistency of resting-state fluctuations across subjects

Our initial analysis suggests that low-frequency resting patterns in fMRI are not only predominantly contained within grey matter, but also appear to be localised within discrete areas of functional significance. In order to investigate this further, we perform an exploratory group analysis for data from 10 subjects obtained under rest (fourth data set).

Data was first motion corrected and co-registered using affine linear registration (Jenkinson and Smith, 2001) into a common space defined by the mid-transformation of all 10 transformation matrices which take the individual low-resolution data sets to the space of the MNI template. This resulted in 10 new EPI data sets which all experienced the same average amount of displacement due to co-registration and which were kept at the original EPI resolution. The individual data sets were then spatially smoothed by a Gaussian kernel with FWHM 7mm and pre-processed by performing voxel-wise variance normalisation and de-meaning as described in section 2. The PCA Eigenbasis, however, was calculated from the mean data covariance matrix, i.e. from the covariance matrix of the spatially concatenated data matrix of size (time \times (voxels \times subject)). All temporal Eigenvectors showed dominant low-frequency content, the initial data reduction stage therefore effectively amounts to a temporal low-pass filtering of the original data by projecting each of the 10 data sets onto the 30 dominant Eigenvectors. The 10 individual (dimensionality-reduced) data sets were then temporally concatenated to form a data matrix of size 300 \times 58032. Using the Laplace approximation to the Bayesian evidence for model order selection, the ICA decomposition estimated 23 spatio-temporal processes. The resulting spatial maps were thresholded using the Gaussian/Gamma alternative hypothesis testing approach. Note that this methodology differs from the Group-ICA methodology introduced by Calhoun et al. (2001) in that the individual data is not projected onto an individual sub-space but instead initial data reduction is performed by projecting each data set onto a common PCA Eigenbasis.

Figure 6 shows example sagittal, coronal and axial slices for 8 spatial patterns (out of 23), overlaid onto the mean subjects' high-resolution structural image (1x1x1.5mm) aligned to the MNI template (all coordinates are in mm from the anterior commissure). The final thresholded maps can be classified as follows:

- **(a) Medial visual cortical areas:** These include primary visual areas located in the calcarine sulcus bilaterally as well as medial, but not lateral, extrastriate regions such as the lingual gyrus. In addition, activation was seen in the inferior division of precuneus cortex and in the lateral geniculate nucleus of the thalamus, the primary 'relay-station' linking visual input to primary visual cortex. Indeed, previous DWI (Diffusion weighted imaging) studies have shown that this region of thalamus connects to the occipital lobe with high probability (Behrens et al. (2003), the thalamic connectivity atlas derived from DWI is available at <http://www.fmrib.ox.ac.uk/connect/>). Although each of these areas (including thalamus) fall within the vascular territory of the posterior cerebral artery, the pattern of results is unlikely to be due to simple temporal patterns of blood supply for two reasons. First, lateral occipito-temporal regions are also supplied by the posterior cerebral artery and these do not exhibit the same resting time course. In fact, one region of the middle temporal gyrus (MT, blue) demonstrated a de-activation at rest. Second, the blood supply to the thalamus is also driven by contributions from other arteries via the circle of Willis, and thus would be unlikely to exhibit the same temporal pattern of blood flow as the posterior cerebral artery alone. Instead, these activations may correspond to a set of connected regions primarily composed of the primary visual system.
- **(b) Lateral visual cortical areas:** These included the occipital pole extending laterally towards the occipito-temporal junction, encompassing non-primary regions of visual cortex. In addition, activation was seen more dorsally in superior parietal regions. This set of regions is frequently found to be co-activated in functional studies of visual attention or visuo-spatial attention. De-activation is found in posterior cingulate cortex (Brodmann's area 30), a region which has been found to co-activate in studies of orientation and navigation in large-scale spaces (Maguire, 2001) and in a more recent study has been found to de-activate in a delayed-match-to-sample task under sleep deprivation (Habeck et al., 2004).
- **(c) Auditory system:** Activation in this component encompassed primary and secondary auditory cortices, including Heschl's gyrus, planum polare and planum temporale, the lateral superior temporal gyrus, and posterior insular cortex (Rivier and Clarke, 1997; Rademacher et al., 2001). Additional activation was observed in the anterior cingulate

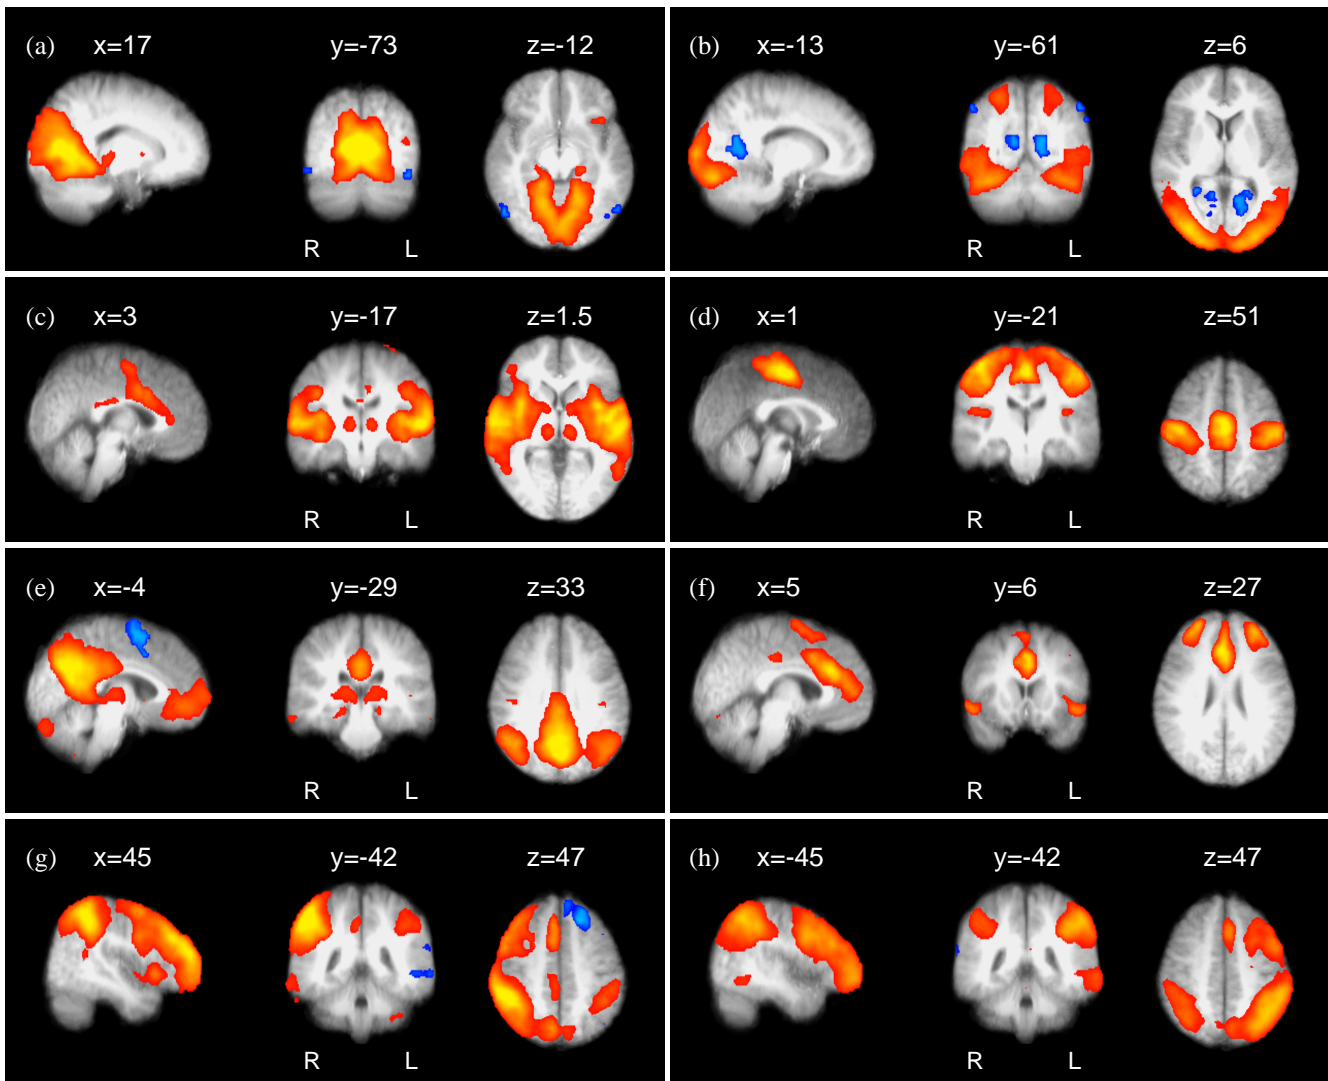


Figure 6: Different PICA-estimated resting patterns: sagittal, coronal and axial views of different spatial maps associated with low-frequency resting patterns estimated from a group of 10 subjects. All images have been co-registered into the space of the MNI template. The coordinates refer to mm distances from the anterior commissure and images are shown in radiological convention.

cortex, anterior supramarginal gyrus, and thalamus. Although these regions are primarily supplied by the middle cerebral artery (MCA), like previous RSNs, this appears to be confined to functionally related areas rather than encompass the entire fronto-parietal area supplied by the MCA.

- **(d) Sensori-motor system:** These included activation in pre- and post-central gyri extending from the superior bank of the Sylvian fissure to the medial wall of the inter-hemispheric fissure and included the supplemental motor area (SMA). This pattern of activation corresponds closely to that seen in bimanual motor tasks. Similar patterns have been identified in previous RSN studies (Biswal et al., 1995).
- **(e) Visuospatial system:** Activation was observed in the posterior parietal cortex at the occipito-parietal junction, along the mid-line in the precuneus and posterior cingulate cortex, and in the frontal pole. Within this spatial map, de-activation was mainly found in pre-SMA. These regions are the same as the physiological baseline areas proposed by Gusnard and Raichle (2001). The authors hypothesise that activity within the posterior cingulate cortex and adjacent precuneus is associated with the representation of the world around us. Similarly, previous studies have demonstrated that lesions affecting the lateral posterior parietal areas lead to severe deficits in spatial attention (Posner et al., 1984; Mesulam, 1981), while disrupting the input to the region prevents patients from benefiting from spatial cues (Alivisatos and Milner, 1989; Koski et al., 1998; Petrides and Pandya, 2002). Similar findings in macaques show that neuronal re-

sponses in a corresponding region (area 7a) are suppressed when stimuli appear at a cued spatial location (Constantinidis and Steinmetz, 2001). Taken together, these findings suggest that the posterior parietal cortex is engaged in orienting to salient visuo-spatial cues. Interestingly, the only region to show deactivation was the pre-SMA, a region consistently implicated in internally (i.e. memory-driven), rather than externally cued, tasks (Rushworth et al., 2004). Overall then, this pattern of findings is consistent with findings by Gusnard and Raichle (2001) that activation in this component may reflect a system of cortical regions engaged in attending to visuo-spatial information. A similar spatial map has been identified using both seed-voxel based correlation analysis and PICA (Greicius et al., 2003, 2004) and has been shown to have reduced activity in the posterior cingulate cortex in Alzheimer’s disease.

- **(f) Executive Control:** Areas of activation include superior and middle prefrontal cortices, anterior cingulate and paracingulate gyri, and ventrolateral prefrontal cortex. In addition, there was sub-cortical activation in a region of the thalamus which connects to pre-frontal cortex with high probability (Behrens et al., 2003; Johansen-Berg et al., 2004). These areas have been hypothesised to provide bias signals to other areas of the brain in order to implement cognitive control (Miller and Cohen, 2001).
- **(g),(h) Dorsal visual stream:** Unlike all of the previous components, these two components showed primarily lateralized activations corresponding roughly to the dorsal visual stream. For instance, (g) shows activation in right lateral occipital complex, right inferior parietal cortex, bilateral intraparietal sulcus, and right middle and superior frontal gyri. The complementary pattern is seen in the left hemisphere in (h), including the bilateral intra-parietal activation. In fact, this is the only area to show a bilateral response in these two components and activation in the region overlaps in the two components, demonstrating that PICA can be used to reveal separate neural systems that contain spatial overlap (i.e. functional integration, cf. Friston (2002)).

5 Discussion

Various researchers have demonstrated that an ICA decomposition can be used to identify patterns of activation, image artefacts and physiologically-generated components including RSNs (De Luca et al., 2002a,b; Kiviniemi et al., 2003; Goldman and Cohen, 2003; Beckmann and Smith, 2004; DeLuca et al., 2004; Greicius et al., 2004). Here we extend the scope of previous investigation into the use of ICA in fMRI, investigating a variety of important aspects relating to the applicability of ICA to resting-state studies. We have demonstrated that an ICA approach can identify a variety of fluctuations (even in cases where these signals coincide at a particular voxel’s location) and that ICA is able to estimate even largely overlapping spatial processes. Furthermore, using high- and low-TR data, we have provided evidence that ICA-derived spatial maps of RSNs are unaffected by respiratory/cardiac fluctuations even though at normal TR the temporal structure of the latter becomes aliased into a frequency range which overlaps that of resting-state fluctuations. Our results suggest that the resting patterns, which qualitatively resemble fMRI activation maps, are largely contained within grey matter and have a different spatial characteristic than ICA maps of major blood vessels. Finally, using data from 10 subjects, we have shown that resting-state patterns are spatially consistent across subjects, clearly identifying networks of functional significance including areas such as visual, sensory or motor cortex as being reproducibly found across subjects.

The above results do not necessarily imply that these spatial patterns are of neuronal origin; they might simply relate to changes in local physiology such as fluctuations linked to local cytoarchitecture and/or local vasculature, which would make RSNs of less interest to neuroscience. As an example, Wise et al. (2004) have demonstrated that large cortical areas of the occipital, parietal and temporal lobes exhibit fluctuations which covary significantly with the end-tidal level of exhaled carbon dioxide. Further studies are required to characterise the extent to which these non-neuronal processes relate to the ICA findings. The maps generated by the ICA group analysis separate these large areas into smaller networks, suggesting that, while there might be a common underlying low-frequency signal induced by vascular processes due to the arterial carbon dioxide fluctuations, these networks have characteristic additional signal fluctuations which can be detected by ICA.

Regardless of their underlying cause, however, RSNs are a major source of structured non-modelled noise in fMRI, and as such deserve to be better understood. Not only do they contribute significantly to the residual variance (lowering sensitivity to true activation), but because they often correlate temporally with experimental paradigm timings, can cause positive or negative bias in the estimated activation (i.e. cause false positive or false negatives).

A limiting factor to the interpretability of PICA-derived maps clearly stems from the fact that the fMRI BOLD signal is an indirect measure of neuronal activity. The spatial sensitivity and specificity from using ICA on resting fMRI data, however, is relatively high, which could be utilised for more advanced approaches, e.g. joint temporal and spatial exploratory data decompositions such as those introduced by Martínez-Montes et al. (2004), which simultaneously explain data from imaging modalities with high temporal (EEG) and high spatial resolution (fMRI).

6 Acknowledgements

The authors wish to thank Heidi Johansen-Berg and Tim Behrens for access to DWI-derived thalamic connectivity data. The authors gratefully acknowledge the financial support from the UK EPSRC, MRC and GlaxoSmithKline.

References

- Alivisatos, B. and Milner, B. (1989). Effects of frontal or temporal lobectomy on the use of advance information in a choice reaction task. *Neuropsychologia*, **27**:495–503.
- Beckmann, C., Noble, J., and Smith, S. (2001). Investigating the intrinsic dimensionality of fMRI data for ICA. In *Seventh Int. Conf. on Functional Mapping of the Human Brain*.
- Beckmann, C. and Smith, S. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans. on Medical Imaging*, **23**(2):137–152.
- Beckmann, C., Woolrich, M., and Smith, S. (2003). Gaussian / Gamma mixture modelling of ICA/GLM spatial maps. In *Ninth Int. Conf. on Functional Mapping of the Human Brain*.
- Behrens, T., Johansen-Berg, H., Woolrich, M., Smith, S., Wheeler-Kingshott, C., Boulby, P., Barker, G., Sillery, E., Sheehan, K., Ciccarelli, O., Thompson, A., Brady, J., and Matthews, P. (2003). Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nature Neuroscience*, **6**(7):750–757.
- Bell, A. and Sejnowski, T. (1995). An information-maximisation approach to blind separation and blind deconvolution. *Neural Computation*, **7**(6):1129–1159.
- Biswal, B., Yetkin, F., V.M., H., and Hyde, J. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, **34**:537–541.
- Bullmore, E., Brammer, M., Williams, S., Rabe-Hesketh, S., Janot, N., David, A., Mellers, J., Howard, R., and Sham, P. (1996). Statistical methods of estimation and inference for functional MR image analysis. *Magnetic Resonance in Medicine*, **35**(2):261–277.
- Calhoun, V., Adali, T., Pearlson, G., and Pekar, J. (2001). A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping*, **14**(3):140–151.
- Comon, P. (1994). Independent component analysis – a new concept? *Signal Processing*, **36**:287–314.
- Constantinidis, C. and Steinmetz, M. (2001). Neuronal responses in area 7a to multiple stimulus displays: Ii. responses are suppressed at the cued location. *Cereb Cortex*, **11**(7):592–597.
- De Luca, M., Beckmann, C., Behrens, T., Clare, S., Matthews, P., De Stefano, N., Woolrich, M., and Smith, S. (2002a). Low frequency signals in fMRI - "resting state networks" and the "intensity normalisation problem". In *Proc. Int. Soc. of Magnetic Resonance in Medicine*.
- De Luca, M., Beckmann, C., Clare, S., Behrens, T., De Stefano, N., Matthews, P., Woolrich, M., and Smith, S. (2002b). Further investigations into "resting state networks" and the "intensity normalisation problem". In *Eighth Int. Conf. on Functional Mapping of the Human Brain*.
- DeLuca, M., Beckmann, C., Clare, S., Matthews, P., De Stefano, N., and Smith, S. (2004). Investigations of resting state networks in fMRI. in preparation.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society Series B*, **39**:1–38.
- Everson, R. and Roberts, S. (2000). Inferring the eigenvalues of covariance matrices from limited, noisy data. *IEEE Transactions on Signal Processing*, **48**(7):2083–2091.
- Friston, K. (1998). Modes or models: A critique on independent component analysis for fMRI. *Trends in Cognitive Sciences*, **2**(10):373–5.
- Friston, K. (2002). Functional integration and inference in the brain. *Prog. Neurobiol.*, **68**(2):113–143.

- Goldman, R. and Cohen, M. (2003). Tomographic distribution of resting alpha rhythm sources revealed by independent component analysis. In *Ninth Int. Conf. on Functional Mapping of the Human Brain*.
- Goldman, R., Stern, J., Engel, J., and M, C. (2002). Simultaneous EEG and fMRI of the alpha rhythm. *NeuroReport*, **18**:2487–92.
- Greicius, M.D., Krasnow, B., Reiss, A.L., and Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *PNAS*, **100**(1):253–258.
- Greicius, M.D., Srivastava, G., Reiss, A.L., and Menon, V. (2004). Default-mode network activity distinguishes alzheimer’s disease from healthy aging: Evidence from functional MRI. *PNAS*, **101**(13):4637–4642.
- Gusnard, D. and Raichle, M. (2001). Searching for a baseline: Functional imaging and the resting human brain. *Nature Review, Neuroscience*, **2**:685–92.
- Habeck, C., Rakitin, B., Moeller, J., Scarmeas, N., Zarahn, E., Brown, T., and Stern, Y. (2004). An event-related fmri study of the neurobehavioral impact of sleep deprivation on performance of a delayed-match-to-sample task. *Cognitive Brain Research*, **18**:306–321.
- Hartvig, N. and Jensen, J. (2000). Spatial mixture modelling of fMRI data. *Human Brain Mapping*, **11**(4):233–248.
- Hyvärinen, A. and Oja, E. (1997). A fast fixed-point algorithm for independent component analysis. *Neural Computation*, **9**(7):1483–1492.
- Jenkinson, M., Bannister, P., Brady, J., and Smith, S. (2002). Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, **17**(2):825–841.
- Jenkinson, M. and Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, **5**(2):143–156.
- Johansen-Berg, H., Behrens, T., Sillery, E., Ciccarelli, O., Thompson, A., S.M., S., and Matthews, P. (2004). Functional-anatomical validation and individual variation of diffusion tractography-based segmentation of the human thalamus. *Cereb Cortex*, in print.
- Johnstone, I. (2000). On the distribution of the largest principal component. Technical report, Department of Statistics, Stanford University.
- Kiviniemi, V., Kantola, J.-H., Jauhiainen, J., Hyvärinen, A., and Tervonen, O. (2003). Independent component analysis of nondeterministic fMRI signal sources. *NeuroImage*, **19**(2):253–260.
- Koski, L., Paus, T., and Petrides, M. (1998). Directed attention after unilateral frontal excisions in humans. *Neuropsychologia*, **36**:1363–1371.
- Krüger, G. and Glover, G.H. (2001). Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magnetic Resonance in Medicine*, **46**:631–637.
- Lowe, M., Mock, B. J., and Sorensen, J.A. (1998). Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *NeuroImage*, **7**(2):119–132.
- Maguire, E. (2001). The retrosplenial contribution to human navigation: a review of lesion and neuroimaging findings. *Scand. J. Psychol.*, **42**(3):225–238.
- Martínez-Montes, E., Valdés-Sosa, P., Miwakeichi, F., Goldman, R., and Cohen, M. (2004). Concurrent eeg/fmri analysis by multi-way partial least squares. *NeuroImage*, **22**(3):1023–1034.
- Mazoyer, B., Zago, L., Bricogne, S., Etard, O., Houde, O., Crivello, F., Joliot, M., Petit, L., and Tzourio-Mazoyer, N. (2001). Cortical networks for working memory and executive functions sustain the conscious resting state in man. *Brain Research Bulletin*, **54**(3):287–98.
- McKeown, M.J., Makeig, S., Brown, G. G., Jung, T. P., Kindermann, S. S., Bell, A. J., and Sejnowski, T. J. (1998). Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping*, **6**(3):160–88.
- Mesulam, M.-M. (1981). A cortical network for directed attention and unilateral neglect. *Annals of Neurology*, **10**:309–325.

- Miller, E. and Cohen, J. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, **24**:167–202.
- Minka, T. (2000). Automatic choice of dimensionality for PCA. Technical Report 514, MIT Media Lab.
- Obrig, H., Neufang, M., Wenzel, R., Kohl, M., Steinbrink, J., Einhaupl, K., and Villringer, A. (2000). Spontaneous low frequency oscillations of cerebral hemodynamics and metabolism in human adults. *NeuroImage*, **12**(6):623–639.
- Petersen, K., Hansen, L., Kolenda, T., Rostrup, E., and Strother, S. (2000). On the independent components of functional neuroimages. In *Second Int. Symp. on Independent Component Analysis and Blind Signal Separation*, 615–620.
- Petrides, M. and Pandya, D. (2002). Association pathways of the prefrontal cortex and functional observations. In Stuss, D. and Knight, R., editors, *Principles of Frontal Lobe Function*, pages 31–50. Oxford University Press.
- Posner, M.I., Walker, J.A., Friedrich, F.J., and Rafal, R.D. (1984). Effects of parietal injury on covert orienting of attention. *Journal of Neuroscience*, **4**:1863–1874.
- Rademacher, J., Morosan, P., Schormann, T., Schleicher, A., Werner, C., Freund, H.-J., and Zilles, K. (2001). Probabilistic Mapping and Volume Measurement of Human Primary Auditory Cortex. *NeuroImage*, **13**(4):669–683.
- Raichle, M., MacLeod, A., Snyder, A., Powers, W., and Gusnard, D. (2001). A default mode of brain function. *PNAS*, **98**(2):676–682.
- Rivier, F. and Clarke, S. (1997). Cytochrome oxidase, acetylcholinesterase, and nadph-diaphorase staining in human supratemporal and insular cortex: Evidence for multiple auditory areas. *NeuroImage*, **6**(4):288–304.
- Roberts, S. and Everson, R., editors (2001). *Independent Component Analysis: Principles and Practice*. Cambridge University Press.
- Rushworth, M., Walton, M., Kennerley, S., and Bannerman, D. (2004). Action sets and decisions in the medial frontal cortex. *Trends Cogn. Sci.*, **8**(9):410–417.
- Shulman, G., Fiez, J., Corbetta, M., Buckner, R., Miezin, F. M., Raichle, M., and Petersen, S. (1997). Common blood flow changes across visual task: Ii. decreases in cerebral cortex. *Journal of Cognitive Neuroscience*, **9**(5):648–63.
- Smith, S.M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, **17**(3):143–155.
- Tipping, M. and Bishop, C. (1999). Mixtures of probabilistic principal component analyzers. *Neural Computation*, **11**(2):443–482.
- Wise, R., Ide, K., Poulin, M.J., and Tracey, I. (2004). Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in bold signal. *NeuroImage*, **21**(4):1652–1664.
- Woolrich, M., Ripley, B.D., Brady, M., and Smith, S.M. (2001). Temporal autocorrelation in univariate linear modelling of fMRI data. *NeuroImage*, **14**(6):1370–1386.
- Xiong, J., Parsons, L.M., Gao, J.H., and Fox, P.T. (1999). Interregional connectivity to primary motor cortex revealed using MRI resting state images. *Human Brain Mapping*, **8**(2-3):151–156.