



Published as: *Curr Opin Neurobiol.* 2003 October ; 13(5): 620–629.

Independent component analysis of functional MRI: what is signal and what is noise?

Martin J McKeown^{*,†}, Lars Kai Hansen[‡], and Terrence J Sejnowski^{§,#}

^{*}Brain Imaging and Analysis Center (BIAC), Duke University, Durham, NC USA

[†]Department of Medicine (Neurology), Duke University, Durham, NC USA

[‡]Informatics and Mathematics Modelling, Technical University of Denmark, B321, DK-2800 Lyngby, Denmark

[§]Howard Hughes Medical Institute, Salk Institute for Biological Studies, 10010 North Torrey Pines Road, La Jolla, CA 92037, USA

[#]Department of Biology, University of California at San Diego, La Jolla, CA 92093, USA

Abstract

Many sources of fluctuation contribute to the functional magnetic resonance imaging (fMRI) signal, complicating attempts to infer those changes that are truly related to brain activation. Unlike methods of analysis of fMRI data that test the time course of each voxel against a hypothesized waveform, data-driven methods, such as independent component analysis and clustering, attempt to find common features within the data. This exploratory approach can be revealing when the brain activation is difficult to predict beforehand, such as with complex stimuli and internal shifts of activation that are not time-locked to an easily specified sensory or motor event. These methods can be further improved by incorporating prior knowledge regarding the temporal and spatial extent of brain activation.

Introduction

In addition to providing a non-invasive, indirect measure of neuronal activity, the blood oxygen level dependent (BOLD) signal in functional magnetic resonance imaging (fMRI) includes contributions from many other sources including the heart beat, breathing and head motion artifacts. There are also less well understood sources, such as low frequency drifts that can also be recorded from cadaver brains [1] and high amplitude oscillations caused by pulse effects that induce localized motion of brain tissue [2].

In contrast to positron emission tomography (PET), in which the measurement represents physiological quantities that can be compared quantitatively with other measurements [3] (e.g. $\mu\text{mol}/100\text{ g tissue}/\text{min}$), the fMRI signals have no simple quantitative physiological interpretation. As a consequence, in most fMRI experiments the signal at a given spatial location (or voxel) during the performance of a task is compared with its value during a period of rest, despite the fact the baseline condition itself contains ongoing cortical activity [4]. Isolating signals of interest is thus a very important problem. In this review we briefly explore traditional and more recently developed methods used to infer task-related changes in fMRI data with a focus on independent component analysis (ICA).

During a neuroimaging experiment, volumetric fMRI signals, acquired as individual slices having a spatial resolution of a few millimeters, are typically sampled with a repetition time (TR) of around 1 Hz. The fMRI signal has temporal and spatial structure at many time and length scales and can be analyzed by different signal processing strategies that emphasize either the spatial or the temporal aspects [5]. One of the most direct ways to estimate whether a given voxel is affected by the behavioral performance or not is to simply cross-correlate the pixel time series with a reference time course describing the sequence of behavioral events. The cross-correlation method can be adapted to take account of the hemodynamic response by first convolving the reference time course with an estimate of the hemodynamic response [6], followed by a voxel-wise t test for significant difference between baseline and activation. Although this method remains popular, its specificity has recently been questioned [7].

Correlation is an example of an hypothesis driven, or confirmatory analysis method, which tests one or more specific hypotheses regarding the time courses of a voxel. By far the most popular software package that uses this approach is statistical parametric mapping (SPM), which employs the general linear model (GLM), an instantiation of multivariate linear regression, and associated methods to deal with violations of the assumptions of the multivariate regression framework, such as the lack of independence among voxels [8]. Studies looking at 'null' datasets, in which a subject does not perform a pre-specified task but merely lies quietly in the scanner, have been valuable in determining the false positive rates that can arise from these approaches [9].

Explorative approaches, which seek to uncover the features of the data themselves, are complementary to hypothesis-driven methods and can help to generate new hypotheses, separate and understand the nature of confounds and find non-trivial components of interest. The main benefit of using a purely data-driven approach to determine the underlying structure of the data is that often the expected time course of brain activation is difficult to specify *a priori*. Two popular data-driven techniques include ICA [10] and different types of clustering in the temporal domain [11–16]. With clustering, a measure is used that estimates the 'similarity' between waveforms, and then all voxels with similar waveforms are collected together within a cluster.

Independent component analysis

The methodology of ICA

Blind signal separation is a class of explorative tools developed for the analysis of images and sound. They are called 'blind' because they aim to recover source signals from mixtures with unknown mixing coefficients. In the cocktail party problem, for example, several microphones in a room record signals from multiple speakers (sources) that arrive with different relative amplitudes at each microphone. ICA is a family of methods for blind signal separation formed on the basis of assumed statistical independence of the source signals [17,18]. The diverse nature of the signals that contribute to fMRI recordings suggests that blind signal separation techniques could be used to isolate these different sources [10,19, 20–22]. Here, we review recent contributions and discuss their results in the context of the basic assumptions of the applied ICA methods.

Let the fMRI signal be represented by the space-time data matrix of measurements $X_{j,t}$, where $j = 1, \dots, J$ (J is the number of pixels/voxels [the three dimensional equivalent of a pixel]); and $t = 1, \dots, T$, (T is the number of time samples). In the linear mixing case we assume that the matrix can be modeled as follows:

$$X_{jt} = \sum_{k=1}^K A_{jk} S_{kt} + E_{jt} \quad (1)$$

where A and S (where the columns of A represent component maps, and the rows of S represent time courses of the respective component maps) are formed by the K independent components of the process, and E is spatially and temporally white noise.

In spatial ICA we assume that the columns of the matrix $A = \lfloor A_{jk} \rfloor$ are statistically independent processes, whereas in temporal ICA the rows of $S = \lfloor S_{jk} \rfloor$ are assumed independent. From the very first application of ICA to fMRI there has been a lively discussion of spatial versus temporal independence, for example see Peterson *et al.*, Friston, Calhoun *et al.* and McKeown *et al.* [23–26]. To appreciate the difference between the two approaches it is interesting to contrast briefly ICA with principal component analysis (PCA).

The basic tool for PCA is singular value decomposition (SVD):

$$X_{jt} = \sum_{k=1}^K U_{jk} \Lambda_{kk} V_{tk} \quad (2)$$

where U and V are orthogonal matrices that are best understood as basis sets that span the spaces of spatial and temporal patterns respectively. The columns of U are the eigenvectors of the Q-mode covariance matrix, which investigates the inter-relationships between voxels: $\Sigma_{jk}^Q = \Sigma_t^T (X_{jt} X_{kt}) / T$, with the columns of V being the eigenvectors of the R-mode covariance matrix, which investigates inter-relationships between volumes at different timepoints: $\Sigma_{st}^R = \Sigma_j^T (X_{js} X_{jt}) / J$. SVD does not allow identification of a mixing matrix. Note, however, that if the model in equation 1 is correct, if the number of sources, K , is small (relative to J and T), and if the variance of the additive noise is small, the important signal variance components will be confined to spatial or temporal subspaces spanned by the K first vectors of either U or V as identified by SVD. Hence, SVD can be used to reduce the dimensionality of the ICA problem [27,28]. PCA and ICA were further compared in the context of denoising by Thomas *et al.* [29].

To completely identify the mixing matrix and the source signals we need to go beyond mere covariance measurements. Two general classes of algorithms use higher order statistics or intra-source correlation. Infomax [30], JADE [31] and FastICA [32] are the most widespread higher order statistics algorithms, whereas the most widespread de-correlation methods are those of Molgedey and Schuster [33] and Ziehe and Muller [34].

The success of a higher order statistics based method depends on how well the source moment structure matches the assumptions of the algorithm, in particular the sign of the fourth central moment, the so-called 'kurtosis'. For the de-correlation based methods, the potential for separation is related to how well separated the source auto-correlation functions are.

When a data analysis problem is approached by ICA there are number of issues to address: what are the independent components in the data? How many components are there? In addition, which ICA algorithm is appropriate?

In many signal processing applications the measurements represent a scene (auditory or visual) or a receptive field in which the independent components naturally reflect independent agents (speakers, objects, mechanical degrees of freedom etc.). When applying

ICA to fMRI the independent source signals are interpreted as networks of similar BOLD activity. In terms of the basic ICA model (equation 1) a single ICA component (say the k th) consists of a spatially distributed set of pixels (A_{jk}) that are activated by the associated time function (S_{kt}). It is useful to visualize the signal reconstructed from one or more components:

$$X_{jt}^{rec} = \sum_{k=1}^K X_{jt}^{(k)}, \quad X_{jt}^{(k)} = \sum_{k=1}^K A_{jk} S_{kt} \quad (3)$$

McKeown *et al.* [27], in the first application of ICA to fMRI, analyzed an fMRI dataset with Infomax, arguing for spatial independence [28]. In this formulation each voxel's time course (row of X in equation 1) is considered a T -dimensional vector, with T being the number of time points in the experiment, and then vectors (time courses) are derived so that the derived time courses (rows of S) have weightings (columns of A) that are as independent as possible. With this application of the algorithm, they found good separation of modes that were task related, transiently task related, as well as confounding modes that represented, for example, head motion [27].

Spatial stationarity, whereby the collection of voxels used in the analysis is assumed to be derived from a single multivariate distribution, is usually assumed by ICA. This may be investigated with test-retest replicability [35•], and possibly addressed with mixture models [36]. In a mixture model, the voxels are partitioned into suitable subsets, and then separate ICA analyses are performed on each subset.

Ordering of components

As is the situation with PCA, we can order the ICA components according to the amount of variance explained:

$$\sigma_k^2 = \sum_{jt} (X_{jt}^{(k)})^2 = \sum_{jt} A_{jk}^2 S_{kt}^2.$$

The total signal variance is approximately the sum of the component variances, hence these variances form a natural ordering. Alternatively, we can order the components according to other features of interest. The most obvious is comparing a component's time course with the behavioral experiment, either by visual inspection [27] or by computing cross-correlations, as mentioned above [23]. In the study by Moritz *et al.* [37••] components are ranked according to frequency content. Among a total of 85 components the power spectrum ranking method identified and ranked the task-related components high, and hence was successful at separating these from artifacts and confounds. Furthermore, ICA was found to be more specific in deriving activated spatial locations than performing power spectrum analysis on the raw time course of each voxel.

A general framework for ordering of components is presented by Lu and Rajapakse [38] enforcing constraints on the Infomax estimation procedure, either of a statistical nature (e.g. ordering components according to kurtosis) or of other types of *a priori* features of interest. In analysis of a visual stimulation fMRI dataset components were sorted according to spatial kurtosis, equivalent to sorting according to spatial sparseness; hence, the most task related component was selected first with high kurtosis, components corresponding to local flow artifacts were also ranked high in the measure. More elaborate and realistic statistical assumptions are invoked in Stone *et al.* and Formisano *et al.* [39,40] for identifying components with asymmetric histograms, autocorrelation or spatial clustering.

Validation of independent component analysis

Temporal versus spatial independent component analysis—New methods for analyzing data, such as ICA, need to be tested on a wide range of problems for robustness and sensitivity to artifacts before the results can be properly interpreted. As there are a variety of different ICA algorithms, it is also important to compare their performance to better understand their strengths and limitations.

Lange *et al.* [5] compared spatial Infomax ICA on simulated and real fMRI data with several other data analysis methods and found that ICA could identify locations of activation not accessible by simple correlation, *t* test or general linear model based methods. The simulated data was created by adding artificial activation foci to real rest fMRI data. The quantitative measure of performance was formed on the basis of receiver operating characteristic (ROC) curves.

Biswal and Ulmer [41] used Infomax ICA to search for temporally independent activation sequences; however, they limited the spatial sample to relatively small regions (30 pixels). Temporal ICA analysis resolved two different induced effects on the fMRI signal: a task induced effect and CO₂ inhalation (hypercapnia). As hypercapnia induces a globally enhanced BOLD signal, searching for spatial independence is not relevant.

A comparative analysis of task related components for a visual stimulation dataset (TR = 0.3 s) using three different spatial and temporal ICA algorithms was presented in Petersen *et al.* [23]. The stimulus reference function was used to identify the consistently task related component in each setup. The consistently task related components showed strong similarities in terms of both the component time series and the spatial maps for all six combinations. The spatial maps derived from the Molgedey-Schuster model were noisier and the associated time courses demonstrated some traces of heartbeat, because for the given task these signal components have similar spatial distributions. The Infomax algorithm worked well when looking for spatially independent patterns such as those in the study by McKeown *et al.* [10]. When the Infomax algorithm looked for temporally independent waveforms, it was less efficient because the boxcar design of the experiment has negative kurtosis. The extended Infomax ICA algorithm can separate components with mixed positive and negative kurtosis and should perform better when looking for temporally independent waveforms amongst voxels [42]. The algorithm developed by Attias [43] that combines higher order statistics and decorrelation works well when looking for temporally and spatially independent patterns, but at considerably higher computational cost. Figure 1 provides an example of ICA components isolated with temporal ICA.

Calhoun *et al.* [44] used the FastICA implementation of ICA and found good correspondence between spatial and temporal modes for an activation study with a single active region; however, for a visual paradigm in which two closely related regions were active they found some divergence between spatial and temporal ICA. The comparison was with the hypothesis driven regression approach using an *a priori* activation time series. Such a comparison is possible for a simple, sustained, sensory stimulus for which the primary activation is predictable. Responses can be less predictable for brief visual stimuli (Duann *et al.* [45••]).

Different independent component analysis algorithms—Esposito *et al.* [46] compared Infomax and FastICA for both simulated and real fMRI data (TR = 3 s) for a standard finger tapping paradigm and visual stimulation following the approach used by Lange *et al.* [5]. Infomax and FastICA produced similar results, as expected as they are formed on the basis of similar statistical assumptions.

The HYBrid ICA (HYBICA) scheme proposed by McKeown [47] invokes a general linear model approach for post-processing of ICA results. This paper also addresses the crucial issue of how many components to keep in the analysis. Similar to the situation with PCA [48] predictive methods can be used, and in the study by McKeown [47] the so-called predicted residual sum of squares (PRESS) statistic was used.

Hoejen-Soerensen *et al.* [49] assumed binary on/off source density and invoked advanced mean-field methods for both spatial and temporal ICA analysis of fMRI. In spatial mode the ICA model was equivalent to time series clustering with multiple concurrent assignments of data. For the mean-field methods it was possible to determine the optimal number of independent components using Bayesian arguments. For the analysis of a visual stimulation dataset the optimal ICs were found to be formed on the basis of representations with less than 10 components.

Reproducibility—Reproducibility, both intra-subject, and inter-subject is a key issue in explorative data analysis, and data-driven methods, given their sensitivity to the underlying structure of the data, may accentuate intersubject variability. Typically fMRI experiments are employed to make statements indicative of a specific population (e.g. all normal subjects greater than age 65). To compare activation across subjects, analysis methods often spatially transform the data. Anatomical images are acquired at the time of the functional images, and a transformation is determined which allows the anatomical images from one subject to be spatially transformed to the anatomical images of another, or all subjects to be transformed to a common exemplar anatomical volume. Applying the same subject-specific spatial transformation to the functional data allows voxel time courses to be directly compared across subjects and tested for task activation. This method may have limited applicability in older populations, however, as they tend to exhibit more spatial variability in fMRI activations possibly as a consequence of compensatory mechanisms [50].

Are components reproducible across subjects? Methods for performing ICA for groups of subjects have been proposed by Calhoun *et al.*, Lukic *et al.* and Svensen *et al.* [51,52,53]. The basic idea is to concatenate the data from several subjects in the spatial dimension, and perform a joint ICA to identify common activation time courses. However, it is unclear whether greater intersubject variability exists in the spatial patterns of activation (in which case finding common ICA time courses across subjects may be suitable) or in the actual time courses of activation, which would support the usual method of spatial transformation. Another way to partially address the problem of intersubject variability is to specify anatomical regions of interest (ROIs), (e.g. 'supplementary motor area') for each subject, a rather laborious process. Presumably ROIs, as opposed to individual voxels that have been spatially transformed to an exemplar volume, would demonstrate less subject-to-subject variability. Covariance between specified ROIs can then be assessed with methods such as structural equation modeling, in which known anatomical connections are used to constrain the model and the strength of connections between regions of interest is estimated [54].

Motion effects revealed with independent component analysis

A consequence of employing a data-driven technique for fMRI data analysis is that it may reveal unpleasant aspects of the data, such as corruption of the data with motion. fMRI data are extremely sensitive to movement, even when it is less than 1 mm, and this may be a limiting factor of the application of this technology in older adults or subjects with brain diseases. As the ICA components are a sensitive reflection of the data, they also tend to be sensitive to all types of movement including abrupt changes and slow, linear drifts [27].

Even data motion-corrected with standard motion correction schemes such as the automated image-registration (AIR) [55] or SPM [8] still produce apparent motion-related ICA

components (Figure 2). This possibly relates to the fact that most motion correction is performed in isolation to the rest of the statistical analysis process, and it is often the endpoint in the analysis pipeline. Motion-correction algorithms typically spatially transform each volume in a time series to an exemplar volume (such as the first volume in the series) by using a measure of similarity [56]. However, standard motion correction schemes seem to minimally affect the predictability of the data (Figure 2), a major component of the performance of statistical models [47,57]. Incorporating potential motion into the ICA framework appears to be a promising approach (Figure 2; [58]).

Some of the sensitivity to movement may be related to the fact that most ICA algorithms are sensitive to outliers, and that voxels at the interface between the brain and other regions such as the skull have markedly different statistical properties. Resampling methods may help assess the robustness of the derived ICA components to outliers [10,59] and robustness itself may be aided by utilizing spatio-temporal *a priori* information [60].

Applications of independent component analysis

Even with the above restrictions in mind, ICA has proved remarkably versatile in several applications in which the brain activation has been hard to predict beforehand. Activity in the visual [61••,62••,63,64••], auditory [65••] and cognitive [66] domains, and even complex social interaction while simultaneously scanning more than one subject [67] have all been investigated with ICA. In a study by Castelo-Branco *et al.* [62••], the data were analyzed with ICA, and spatial components which loaded heavily on the motion-sensitive visual area hMT(+)/V5 were further examined to determine a functionally connected network involved in perceptual decision. Being able to meaningfully interpret fMRI experiments incorporating very complex stimuli, such as driving [64••] or watching a movie [61••], exemplifies the advantages of not having to specify activation profiles beforehand. ICA has also provided insight into artifacts caused by large vessel effects in perfusion imaging [68]. The versatility of the technique is such that ICA has been used as a preprocessing method for attempting to visualize color multichannel magnetic resonance data [69], and has been considered within the data mining framework used in statistics [70].

In the study by Zeki *et al.* [61••] an experiment is set up to explore the role of the kinetic orbital area (KO) in humans, which is activated when we perceive shapes generated by kinetic boundaries, for example in random dot patterns. ICA analysis was performed on fMRI data from eight subjects while they watched 20 minutes of an action movie. Regions were related if they appeared active in the same spatial component. Using this spatial grouping approach it was found that the KO area was less specific to kinetic boundary activation but activated for stimuli similar to those that activate V3, an area that represents depth and contours.

Using a pluralistic approach involving both ICA and general linear models Calhoun *et al.* [25] investigated a visual perception task that was designed to provide a reliable and valid measure of visual perceptual capacity. The general linear model, which measures the pixel wise correlation with the stimulus reference function, and the ICA yielded similar but not identical results, and together suggested a significant role for the cerebellum in visual processing. The general linear model appeared more selective and sensitive to primary visual and cerebellar regions. However, the ICA also detected primary motor activity, whereas the general linear model did not, the main reason being that the activation time course showed a considerably longer duration than expected on the basis of the standard hemodynamic response model that is invoked in the general linear model estimate. This is a prototypical situation in which the explorative nature of ICA allows detection of an unexpected response.

Using ICA, Duann *et al.* [45••] demonstrated that the common assumption of a spatially and temporally stationary hemodynamic response function does not hold. Subjects exhibited a wide variety of responses to short stimuli including responses with two positive peaks.

The ability of ICA to detect transient and randomly occurring neuropsychological events was studied by Gu *et al.* [66]. An auditory sentence-monitoring fMRI dataset was analyzed and components post-processed to favor components with spatial connectivity. The ICA time courses were shown to match well with button-press signals used to monitor subjects listening to randomized auditory stimuli.

The complex dynamics of neural activation during simulated driving were investigated with ICA by Calhoun *et al.* [64••]. The differential response of several systems including error monitoring, motor control and vigilance could be quantified in more detail using ICA than when using simple subtractive analyses.

If confounding components are excluded in the reconstruction of equation 3, ICA can be viewed as a denoising filter similar to PCA and other signal processing subspace methods. This approach was pursued for electroencephalographic (EEG) analysis by Jung *et al.* [19•] and for fMRI analysis by Thomas *et al.* [29•]. Arfanakis *et al.* [71] pursued a complementary approach and reconstructed a signal matrix without the task induced components to study low frequency oscillations. This was previously done using fMRI resting state data, however, Arfanakis *et al.* argue that as 'rest' data contains ongoing cortical activity, it should be possible to use ICA to remove the activation in a situation in which the brain is focused on a particular task and hence, less subject to random activations [71]. ICA turned out to be very efficient for removing task related parts of the data and the subsequent analysis of the low frequency oscillations showed strong similarities with those observed in the resting state. Although the authors suggest that they have isolated regions demonstrating 'functional connectivity', it is not clear what the underlying mechanism is, as a physiological origin such as the vasculature cannot be ruled out [72].

Other data-driven methods

Rather than using ICA or clustering, a more direct way to determine the time course of response is to use event-related designs [73]. This paradigm is adapted from the EEG and event-related potentials (ERPs) literature, whereby many similar stimuli are presented and the time course (of in this case, a voxel) within a specified time window is averaged, time-locked to stimulus presentation. Unlike the EEG with its excellent temporal resolution, the sluggishness of the hemodynamic response in fMRI (lasting several seconds in duration) places a limit on the speed that stimuli can be presented at. In order that event-related fMRI experiments not become excessively long, the total number of stimuli is restricted to a hundred or so. Standard analyses of event related designs make two key assumptions, first, that the brain response to the stimulus is independent of the brain state, and second, that there is minimal temporal 'jitter' in the fMRI response after presentation of a given stimulus. The trial-to-trial variability has recently been shown to be fairly significant [45••,74], and the implications this has on the overall interpretation of activation in event-related fMRI studies is unclear. Whether or not the other structured (and possibly non-brain) signals will tend to zero when averaged over the many fewer trials than those typically used in ERPs may also need to be more fully investigated [74].

Still, other data-driven analysis methods have attempted to isolate task-related signals from other sources of variability within fMRI data. With a canonical correlation analysis (CCA) approach [75] components in the data with time courses having a broad autocorrelation can be extracted. This method will be robust to transient noise signals caused by abrupt movement, and will tend to isolate slowly varying components, such as those related to the

task but also other sources of slow drifts in the signal not related to brain activation. Investigation of the frequency spectra at each voxel with sophisticated techniques such as multi-taper analysis has allowed the isolation of small but significant frequencies in voxels' time series that are worthy of further investigation [76].

Perhaps reassuringly, data-driven methods like ICA often give comparable responses to traditional hypothesis-based approaches [44], and in some cases with incorrect task performance ICA appeared to provide more accurate maps [77••]. Methods that attempt to combine the strengths of complementary analysis approaches may prove a powerful tool. ICA components can be used to make reasonable determinations of task-related regressors in a general linear model framework [47], or ICA may be used to remove the confound of task-related activation in exploring functional connectivity [78].

However, despite the potential advantages of data-driven methods, the nature of brain activation is inherently a spatio-temporal process. There may be advantages to investigating the spatial patterns associated with activation, for example using spatial Gaussian mixture models [79]. fMRI investigators do have some prior biases as to what constitutes fMRI activation; we tend to distrust isolated single 'activated' voxels as false positives. A grafting of current data driven methods with regularized spatio-temporal solutions [80••] may prove a powerful means to more accurately isolate presumed brain activity.

Conclusions

ICA is a promising exploratory technique that provides an alternative means to view data and to test assumptions about traditional hypothesis-driven methods. Much technical effort has been put into tests to ensure robust inferences about brain activity. A significant virtue of ICA is that it allows the detection of unexpected responses to stimuli, including random responses or transiently task related responses. Furthermore, ICA is an effective tool for denoising fMRI, both with respect to random noise and confounding signals such as pulsation and breathing artifacts. Such techniques will allow the full spatial-temporal aspects of brain activation to be better isolated from the complex mixtures of (often unknown) sources that make up the measured fMRI signal.

Acknowledgments

This survey was supported by a National Institute of Neurologic diseases and Stroke grant P01NS41328-01A1 (core C) to M McKeown, a National Institute of Biomedical Imaging and Bioengineering grant P20 EB 002013 to L Hansen, and both the Howard Hughes Medical Institute and a National Institutes of Health grant MH61619-03 to T Sejnowski.

Abbreviations

BOLD	blood oxygenation level dependent
ERP	event related potentials
fMRI	functional magnetic resonance imaging
ICA	independent component analysis
KO	kinetic orbital area
PCA	principal component analysis
ROI	regions of interest
SPM	statistical parametric mapping

SVD	singular value decomposition
TR	repetition time

References and recommended reading

Papers of particular interest, published within the annual period of review, have been highlighted as:

- of special interest
- of outstanding interest

1. Smith AM, Lewis BK, Ruttimann UE, Ye FQ, Sinnwell TM, Yang Y, Duyn JH, Frank JA. Investigation of low frequency drift in fMRI signal. *Neuroimage*. 1999; 9:526–533. [PubMed: 10329292]
2. Jiang H, Golay X, van Zijl PC, Mori S. Origin and minimization of residual motion-related artifacts in navigator-corrected segmented diffusion-weighted EPI of the human brain. *Magn Reson Med*. 2002; 47:818–822. [PubMed: 11948746]
3. Buchanan, J. Principles and Practice of Positron Emission Tomography. Wahl, R., editor. Williams & Wilkins Publishers; Lippincott: 2002.
4. Gusnard DA, Raichle ME. Searching for a baseline: functional imaging and the resting human brain. *Nat Rev Neurosci*. 2001; 2:685–694. [PubMed: 11584306]
5. Lange N, Strother SC, Anderson JR, Nielsen FA, Holmes AP, Kolenda T, Savoy R, Hansen LK. Plurality and resemblance in fMRI data analysis. *Neuroimage*. 1999; 10:282–303. [PubMed: 10458943]
6. Bandettini PA, Jesmanowicz A, Wong EC, Hyde JS. Processing strategies for time-course data sets in functional MRI of the human brain. *Magn Reson Med*. 1993; 30:161–173. [PubMed: 8366797]
7. Baumgartner R, Somorjai R, Summers R, Richter W, Ryner L. Correlator beware: correlation has limited selectivity for fMRI data analysis. *Neuroimage*. 2000; 12:240–243. [PubMed: 10913329]
8. Friston, KJ. Statistical parametric mapping and other analyses of functional imaging data. In: Toga, AW.; Mazziotta, JC., editors. *Brain Mapping, The Methods*. edn1. Academic Press; 1996. p. 363-396.
9. Zarahn E, Aguirre GK, D'Esposito M. Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage*. 1997; 5:179–197. [PubMed: 9345548]
10. McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Sejnowski TJ. Analysis of fMRI data by blind separation into independent spatial components. *Hum Brain Mapp*. 1998; 6:160–188. [PubMed: 9673671]
11. Filzmoser P, Baumgartner R, Moser E. A hierarchical clustering method for analyzing functional MR images. *Magn Reson Imaging*. 1999; 17:817–826. [PubMed: 10402588]
12. Baune A, Sommer FT, Erb M, Wildgruber D, Kardatzki B, Palm G, Grodd W. Dynamical cluster analysis of cortical fMRI activation. *Neuroimage*. 1999; 9:477–489. [PubMed: 10329287]
13. Golay X, Kollias S, Stoll G, Meier D, Valavanis A, Boesiger P. A new correlation-based fuzzy logic clustering algorithm for fMRI. *Magn Reson Med*. 1998; 40:249–260. [PubMed: 9702707]
14. Moser E, Diemling M, Baumgartner R. Fuzzy clustering of gradient-echo functional MRI in the human visual cortex. Part II: quantification. *J Magn Reson Imaging*. 1997; 7:1102–1108. [PubMed: 9400855]
15. Cordes D, Haughton V, Carew JD, Arfanakis K, Maravilla K. Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magn Reson Imaging*. 2002; 20:305–317. [PubMed: 12165349]
16. Yee SH, Gao JH. Improved detection of time windows of brain responses in fMRI using modified temporal clustering analysis. *Magn Reson Imaging*. 2002; 20:17–26. [PubMed: 11973026]

17. Brown GD, Yamada S, Sejnowski TJ. Independent component analysis at the neural cocktail party. *Trends Neurosci.* 2001; 24:54–63. [PubMed: 11163888]
18. Stone JV. Independent component analysis: an introduction. *Trends Cogn Sci.* 2002; 6:59–64. [PubMed: 15866182]
19. Jung TP, Makeig S, McKeown MJ, Bell AJ, Lee TW, Sejnowski TJ. Imaging brain dynamics using independent component analysis. *Proc IEEE.* 2001; 89:1107–1122. The authors present a comprehensive and critical review of ICA applied to natural images, EEG, ERP, and fMRI.
20. Lee TW, Girolami M, Sejnowski TJ. Independent component analysis using an extended infomax algorithm for mixed subgaussian and supergaussian sources. *Neural Comput.* 1999; 11:417–441. [PubMed: 9950738]
21. Hyvarinen, A.; Karhunen, J.; Oja, E. *Independent Component Analysis.* John Wiley & Sons; 2001.
22. Calhoun, V.; Adali, T.; Hansen, LK.; Larsen, J.; Pekar, J. ICA of Functional MRI Data: an Overview. 4th International Symposium ICA2003; 2003. p. 281-288.
23. Petersen, K.; Hansen, L.; Koleda, T.; Rostrup, E.; Strother, S. On the independent components of functional neuroimages. Third International Conference on Independent Component Analysis and Blind Source Separation (ICA2000); 2000. p. 615-620.
24. Friston KJ. Modes or models: a critique on independent component analysis for fMRI. *Trends Cogn Sci.* 1998; 2:373–375. [PubMed: 21227247]
25. Calhoun VD, Adali T, Pearlson GD, Pekar JJ. Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Hum Brain Mapp.* 2001; 13:43–53. [PubMed: 11284046]
26. McKeown MJ, Makeig S, Brown G, Kindermann S, Sejnowski TJ. Modes or models: a critique on independent component analysis for fMRI - reply to commentary by Prof Friston. *Trends Cogn Sci.* 1998; 2:375. [PubMed: 21227248]
27. McKeown MJ, Jung TP, Makeig S, Brown G, Kindermann SS, Lee TW, Sejnowski TJ. Spatially independent activity patterns in functional MRI data during the stroop color-naming task. *Proc Natl Acad Sci USA.* 1998; 95:803–810. [PubMed: 9448244]
28. Hansen, LK. Blind separation of noisy image mixtures. In: Girolami, M., editor. *Advances in Independent Component Analysis, Perspectives in Neural Computing.* Springer-Verlag; 2000. p. 165-187.
29. Thomas CG, Harshman RA, Menon RS. Noise reduction in BOLD-based fMRI using component analysis. *Neuroimage.* 2002; 17:1521–1537. [PubMed: 12414291] This study uses PCA and ICA to 'de-noise' fMRI data. After artifact and unstructured noise removal the contrast sensitivity was increased. Although PCA is superior for removal of unstructured noise ICA is recommended for removal of structured noise (artifacts).
30. Bell AJ, Sejnowski TJ. An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.* 1995; 7:1129–1159. [PubMed: 7584893]
31. Cardoso J, Souloumiac A. Blind beamforming for non-Gaussian signals. *IEEE Proceedings-F.* 1993; 140:362–370.
32. Hyvarinen A. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans Neural Netw.* 1999; 10:626–634. [PubMed: 18252563]
33. Molgedey L, Schuster H. Separation of independent signals using time-delayed correlations. *Phys Rev Lett.* 1994; 72:3634–3637. [PubMed: 10056251]
34. Ziehe, A.; Muller, K-R. TDSEP - an efficient algorithm for blind separation using time structure. In: Niklasson, L.; Boden, M.; Ziemke, T., editors. 8th ICANN, *Perspectives in Neural Computing;* Berlin. Berlin: Springer Verlag; 1998. p. 675-680.
35. Nybakken GE, Quigley MA, Moritz CH, Cordes D, Haughton VM, Meyerand ME. Test-retest precision of functional magnetic resonance imaging processed with independent component analysis. *Neuroradiology.* 2002; 44:403–406. [PubMed: 12012124] A study demonstrating that the test-retest precision with ICA is similar to that of standard *t* tests.
36. Lee T, Lewicki M, Sejnowski T. ICA mixture models for unsupervised classification of non-gaussian sources and automatic context switching in blind signal separation. *IEEE Transactions on PAMI.* 2000; 22:1–12.

- 37•• Moritz CH, Rogers BP, Meyerand ME. Power spectrum ranked independent component analysis of a periodic fMRI complex motor paradigm. *Hum Brain Mapp.* 2003; 18:111–122. [PubMed: 12518291] In this study independent components are ranked according to frequency content. Among a total of 85 components the power spectrum ranking method identified and ranked the task-related components high, hence separating these from artifacts and confounds. Furthermore, ICA was found to be more specific to derive activated spatial locations than voxel wise power spectrum analysis.
38. Lu W, Rajapakse J. Eliminating indeterminacy in ICA. *Neurocomputing.* 2003; 50:271–290.
39. Stone JV, Porrill J, Porter NR, Wilkinson ID. Spatiotemporal independent component analysis of event-related fMRI data using skewed probability density functions. *Neuroimage.* 2002; 15:407–421. [PubMed: 11798275]
40. Formisano E, Esposito F, Kriegeskorte N, Tedeschi G, Di Salle F, Goebel R. Spatial independent component analysis of functional magnetic resonance imaging time-series: characterization of the cortical components. *Neurocomputing.* 2002; 49:241–254.
41. Biswal BB, Ulmer JL. Blind source separation of multiple signal sources of fMRI data sets using independent component analysis. *J Comput Assist Tomogr.* 1999; 23:265–271. [PubMed: 10096335]
42. Lee TW, Girolami M, Sejnowski TJ. Independent component analysis using an extended infomax algorithm for mixed subgaussian and supergaussian sources. *Neural Comput.* 1999; 11:417–441. [PubMed: 9950738]
43. Attias H, Schreiner CE. Blind source separation and deconvolution by dynamic component analysis. *Neural Networks for Signal Processing VII: Proceedings of the 1997 IEEE Workshop.* 1998; 10:1373–1424.
44. Calhoun VD, Adali T, McGinty VB, Pekar JJ, Watson TD, Pearlson GD. fMRI activation in a visual-perception task: network of areas detected using the general linear model and independent components analysis. *Neuroimage.* 2001; 14:1080–1088. [PubMed: 11697939]
- 45•• Duann JR, Jung TP, Kuo WJ, Yeh TC, Makeig S, Hsieh JC, Sejnowski TJ. Single-trial variability in event-related BOLD signals. *Neuroimage.* 2002; 15:823–835. [PubMed: 11906223] Although a constant hemodynamic response function is a premise for many hypothesis testing schemes ICA is used in this study to discuss the sources of hemodynamic variance. A visual stimulus was administered at 8 Hz in 0.1 sec and 3 sec blocks. fMRI data (TR = 0.5 sec) demonstrated that component hemodynamic responses varied substantially across sessions, subjects, and brain areas. In four of six subjects the response even showed multiple 'peaks'. The underlying cause of this behavior is unknown and the complex response to short stimuli is speculated to reflect top-down attentional modulation.
46. Esposito F, Formisano E, Seifritz E, Goebel R, Morrone R, Tedeschi G, Di Salle F. Spatial independent component analysis of functional MRI time-series: to what extent do results depend on the algorithm used? *Hum Brain Mapp.* 2002; 16:146–157. [PubMed: 12112768]
47. McKeown MJ. Detection of consistently task-related activations in fMRI data with HYBrid independent component analysis (HYBICA). *Neuroimage.* 2000; 11:24–35. [PubMed: 10686114]
48. Hansen LK, Larsen J, Nielsen FA, Strother SC, Rostrup E, Savoy R, Lange N, Sidtis N, Svarer C, Paulson OB. Generalizable patterns in neuroimaging: how many principal components. *Neuroimage.* 1999; 9:534–544. [PubMed: 10329293]
49. Hoejen-Soerensen P, Winther O, Hansen LK. Analysis of functional neuroimages using ICA with adaptive binary sources. *Neurocomputing.* 2002; 49:213–225.
50. Nielson KA, Langenecker SA, Garavan H. Differences in the functional neuroanatomy of inhibitory control across the adult life span. *Psychol Aging.* 2002; 17:56–71. [PubMed: 11931287]
51. Calhoun VD, Adali T, Pearlson GD, Pekar JJ. A method for making group inferences from functional MRI data using independent component analysis. *Hum Brain Mapp.* 2001; 14:140–151. [PubMed: 11559959]
- 52•• Lukic AS, Wemick MN, Hansen LK, Anderson J, Strother SC. A spatially robust ICA algorithm for multiple fMRI data sets. *IEEE International Symposium on Biomedical Imaging.* 2002:839–842. In this study multiple subjects' fMRI scans are concatenated in the spatial dimension to search for common and individual temporal effects. The Molgedey-Schuster algorithm is derived

in terms of prediction matrices and is modified to explicitly handle the structure for common and individual temporal effects.

53. Svensen M, Kruggel F, Benali H. ICA of fMRI group study data. *Neuroimage*. 2002; 16:551–563. [PubMed: 12169242]
54. McIntosh AR, Grady CL, Haxby JV, Ungerleider LG, Horwitz B. Changes in limbic and prefrontal functional interactions in a working memory task for faces. *Cereb Cortex*. 1996; 6:571–584. [PubMed: 8670683]
55. Woods RP, Grafton ST, Holmes CJ, Cherry SR, Mazziotta JC. Automated image registration: I. General methods and intrasubject, intramodality validation. *J Comput Assist Tomogr*. 1998; 22:139–152. [PubMed: 9448779]
56. Jenkinson M, Bannister P, Brady M, Smith S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*. 2002; 17:825–841. [PubMed: 12377157]
57. Myers, RH. *Classical and Modern Regression with Applications*. edn 1. PWS Publishers; Boston, MA: 1986.
58. Liao, R.; McKeown, M.; Krolak, J. 2002 IEEE International Symposium on Biomedical Imaging. Washington, DC: 2002. Movement correction of FMRI time-series using intrinsic statistical properties of images: an independent component analysis approach; p. 765-768.
59. LaConte S, Anderson J, Muley S, Ashe J, Frutiger S, Rehm K, Hansen LK, Yacoub E, Hu X, Rottenberg D, et al. The evaluation of preprocessing choices in single-subject BOLD fMRI using NPAIRS performance metrics. *Neuroimage*. 2003; 18:10–27. [PubMed: 12507440]
60. Suzuki K, Kiryu T, Nakada T. Fast and precise independent component analysis for high field fMRI time series tailored using prior information on spatiotemporal structure. *Hum Brain Mapp*. 2002; 15:54–66. [PubMed: 11747100] By incorporating prior temporal information and suitable contrast functions the authors make a nice attempt to derive a tailored ICA-like algorithm perhaps more suited for fMRI.
61. Zeki S, Perry RJ, Bartels A. The processing of kinetic contours in the brain. *Cereb Cortex*. 2003; 13:189–202. [PubMed: 12507950] This study investigates the role of the so-called kinetic orbital (KO) area. Subjects were fMRI scanned while watching 22 minutes of an action movie. From a total of 324 independent components a subset was selected for voxels in KO and two other visual areas. The time courses of the independent components (ICs) from the KO area were then compared with the ICs from V5 and V3v. On the basis of the correlation coefficient the authors conclude that KO is more akin to V3 than V5, and hence, less specific to processing of kinetic contours than previously believed.
62. Castelo-Branco M, Formisano E, Backes W, Zanella F, Neuenschwander S, Singer W, Goebel R. Activity patterns in human motion-sensitive areas depend on the interpretation of global motion. *Proc Natl Acad Sci USA*. 2002; 99:13914–13919. [PubMed: 12368476] This study demonstrates the benefits of both data-driven and hypothesis-driven methods. The authors present overlapping visual gratings to subjects that could be interpreted as two separate objects or perceived as a single intermediate object. By investigating the ICA component that loaded heavily in region V5, they isolated a distributed brain region whose time course almost matched that of the subjects' perception.
63. Bartels A, Zeki S. The architecture of the colour centre in the human visual brain: new results and a review. *Eur J Neurosci*. 2000; 12:172–193. [PubMed: 10651872]
64. Calhoun VD, Pekar JJ, McGinty VB, Adali T, Watson TD, Pearlson GD. Different activation dynamics in multiple neural systems during simulated driving. *Hum Brain Mapp*. 2002; 16:158–167. [PubMed: 12112769] In this study fMRI data were acquired while subjects were fixating, watching a simulated driving scene, and actively simulated driving in a commercial game simulator (TR = 1 sec). Five consistent and one transiently activated component out of 25 were found. Consistently activated components reflected differential activity in response to running speed in the simulation. A conventional subtractive analysis found the relevant areas of the consistently activated components but mixed the independent components and furthermore missed the transiently activated component associated with task switching.
65. Seifritz E, Esposito F, Hennel F, Mustovic H, Neuhoff JG, Bilecen D, Tedeschi G, Scheffler K, Di Salle F. Spatiotemporal pattern of neural processing in the human auditory cortex. *Science*.

2002; 297:1706–1708. [PubMed: 12215648] In this study the spatio-temporal structure of the human auditory cortex is studied by a combination of spatial and temporal ICA. First spatial ICA is performed and used to isolate a set of regions of interest. The associated time course showed complex behavior and the component was subsequently decomposed by temporal ICA to produce consistent and transient components of activity with maps showing several overlapping spatial compartments of the auditory cortex.

66. Gu H, Engelen W, Feng H, Silbersweig DA, Stern E, Yang Y. Mapping transient, randomly occurring neuropsychological events using independent component analysis. *Neuroimage*. 2001; 14:1432–1443. [PubMed: 11707099]
67. Montague PR, Berns GS, Cohen JD, McClure SM, Pagnoni G, Dhamala M, Wiest MC, Karpov I, King RD, Apple N, et al. Hyperscanning: simultaneous fMRI during linked social interactions. *Neuroimage*. 2002; 16:1159–1164. [PubMed: 12202103]
68. Carroll TJ, Haughton VM, Rowley HA, Cordes D. Confounding effect of large vessels on MR perfusion images analyzed with independent component analysis. *AJNR Am J Neuroradiol*. 2002; 23:1007–1012. [PubMed: 12063234]
69. Muraki S, Nakai T, Kita Y, Tsuda K. An attempt for coloring multichannel MR imaging data. *IEEE Transactions on Visualization and Computer Graphics*. 2001; 7:265–274.
70. Megalooikonomou V, Ford J, Shen L, Makedon F, Saykin A. Data mining in brain imaging. *Stat Methods Med Res*. 2000; 9:359–394. [PubMed: 11084713]
71. Arfanakis K, Cordes D, Haughton VM, Moritz CH, Quigley MA, Meyerand ME. Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. *Magn Reson Imaging*. 2000; 18:921–930. [PubMed: 11121694]
72. Mitra PP, Ogawa S, Hu X, Ugurbil K. The nature of spatiotemporal changes in cerebral hemodynamics as manifested in functional magnetic resonance imaging. *Magn Reson Med*. 1997; 37:511–518. [PubMed: 9094072]
73. Buckner RL. Event-related fMRI and the hemodynamic response. *Hum Brain Mapp*. 1998; 6:373–377. [PubMed: 9788075]
74. McKeown MJ, Varadarajan V, Huettel S, McCarthy G. Deterministic and stochastic features of fMRI data: implications for analysis of event-related experiments. *J Neurosci Methods*. 2002; 118:103–113. [PubMed: 12204302]
75. Friman O, Borga M, Lundberg P, Knutsson H. Exploratory fMRI analysis by autocorrelation maximization. *Neuroimage*. 2002; 16:454–464. [PubMed: 12030831]
76. Mitra PP, Pesaran B. Analysis of dynamic brain imaging data. *Biophys J*. 1999; 76:691–708. [PubMed: 9929474]
- 77••. Quigley MA, Haughton VM, Carew J, Cordes D, Moritz CH, Meyerand ME. Comparison of independent component analysis and conventional hypothesis-driven analysis for clinical functional MR image processing. *AJNR Am J Neuroradiol*. 2002; 23:49–58. [PubMed: 11827875] The authors present a comparison of ICA and regression approaches to explore the similarities and differences in the spatial maps of activation derived by the two methods.
78. Arfanakis K, Cordes D, Haughton VM, Moritz CH, Quigley MA, Meyerand ME. Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. *Magn Reson Imaging*. 2000; 18:921–930. [PubMed: 11121694]
79. Hartvig NV, Jensen JL. Spatial mixture modeling of fMRI data. *Hum Brain Mapp*. 2000; 11:233–248. [PubMed: 11144753]
- 80••. Purdon PL, Solo V, Weisskoff RM, Brown EN. Locally regularized spatiotemporal modeling and model comparison for functional MRI. *Neuroimage*. 2001; 14:912–923. [PubMed: 11554810] The authors make the first step in imposing regularized spatiotemporal solutions on fMRI analysis. As we have some prior knowledge on the spatial extent and temporal dynamics of brain activity, explicitly incorporating this into the solutions may result in more accurate estimates of brain activation.

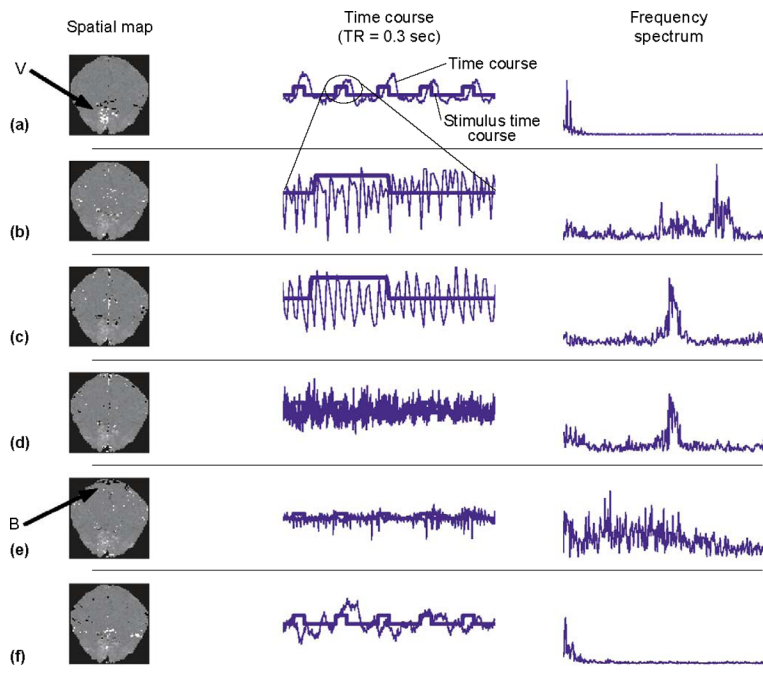


Figure 1.

Results of applying temporal ICA to single-slice fMRI data. The subject was shown a flashing (8 Hz) annular checkerboard pattern interleaved with periods of fixation. There were five runs of 30 scans of fixation (10.0 s), 31 scans of stimulation (10.3 s), and 60 scans of post-stimulus fixation (20.0 s). The power spectrum is estimated in the range 0–1.5 Hz (Nyquist frequency). The slice is aligned with the calcarine sulcus and contains a portion of the primary visual areas. The six independent components shown are represented by the spatial map (the 2.5% highest and lowest values are shown as white and black pixels on a background formed by the average of the dataset providing anatomical references). The components are sorted according to variance contribution. **(a)** The first IC loads heavily in primary visual areas (V) (*left column*), and its time course (*middle column, thin line*) closely follows the stimuli time course (*middle column, thick line*). The power spectrum (*right column, thick line*) of the time course and stimulus time course are closely matched. **(b)** The second component contains pulsations related to the heartbeat as demonstrated in the time course and power spectrum. **(c,d)** The third and fourth components appear related to slower breathing-related periodic confounds. **(e)** Component five is a white noise (broad band) component with a more spiky character, and the component image is dominated by the (negative) boundary area (B), suggesting that this is mostly related to motion artifact. **(f)** The sixth element is a low-frequency component with a period of about 10–15 s unrelated to the stimulus sequence and possibly represents an artifact related to vasomotor oscillations [72].

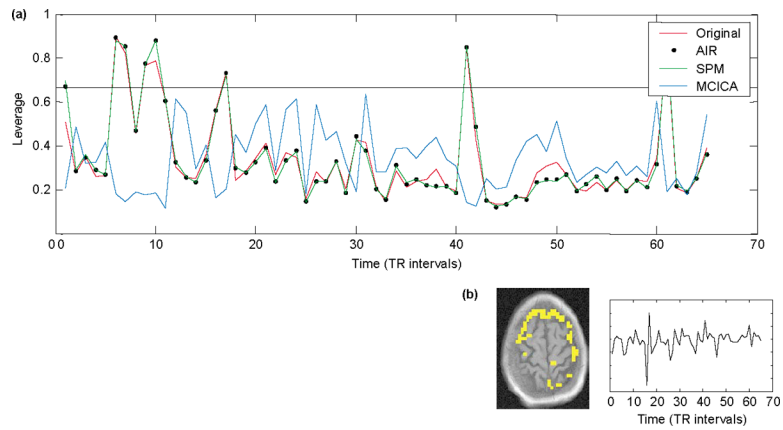


Figure 2. Effects of motion correction on ICA components. **(a)** The predictability of the data, estimated by the diagonals of the Hat matrix, $H = X(X'X)^{-1}X'$, where the columns of X represent the largest 1/3 of the eigenvectors of the covariance matrix, is plotted. The horizontal line is a heuristic used in regression to imply high leverage points. Note that common motion correction schemes (AIR and SPM) do not measurably affect predictability. MCICA = motion corrected ICA. **(b)** Even after standard motion correction, ICA components indicative of movement can still be isolated.