NeuroImage 62 (2012) 891–901

Contents lists available at SciVerse ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

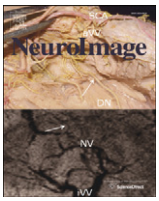

Review

Modelling with independent components

Christian F. Beckmann *

MIRA Institute for Biomedical Technology and Technical Medicine, University of Twente, Enschede, The Netherlands
Donders Institute, Centre for Cognitive Neuroimaging, Radboud University Nijmegen, Nijmegen, The Netherlands
Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB), Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, UK

ARTICLE INFO

Article history:
Accepted 9 February 2012
Available online 18 February 2012

Keywords:
Independent Component Analysis
ICA
Functional Magnetic Resonance Imaging
FMRI
Exploratory data analysis
Factor analysis

ABSTRACT

Independent Component Analysis (ICA) is a computational technique for identifying hidden statistically independent sources from multivariate data. In its basic form, ICA decomposes a 2D data matrix (e.g. time × voxels) into separate components that have distinct characteristics. In FMRI it is used to identify hidden FMRI signals (such as activations). Since the first application of ICA to Functional Magnetic Resonance Imaging (FMRI) in 1998, this technique has developed into a powerful tool for data exploration in cognitive and clinical neurosciences. In this contribution to the commemorative issue *20 years of FMRI* I will briefly describe the basic principles behind ICA, discuss the probabilistic extension to ICA and touch on what I think are some of the most notorious loose ends. Further, I will describe some of the most powerful 'killer' applications and finally share some thoughts on where I believe the most promising future developments will lie.

Contents

Introduction

An increasing number of areas that employ statistical data analysis techniques for scientific investigation operate on data that has been generated from underlying signals of interest by means of complicated, and very often poorly understood, processes. This is certainly the case for Functional Magnetic Resonance Imaging. Here, brain activation at the neuronal level exhibits itself via the BOLD-response to stimulation. The rather poor signal-to-noise ratio suggests that this signal is further obscured by various other sources of variability, possibly including machine artefacts, physiological pulsation, head motion and haemodynamic changes induced by different processes (Toga and Mazziotta, 2002). This mixture of signals presents a huge challenge for analytical methods attempting to identify signals of interest. Instead of operating on data that directly reflects the object of interest, data analysis has to proceed on indirect measurements

* Donders Institute, Centre for Cognitive Neuroimaging (DCCN), Radboud University, P.O. Box 9101, NL-6500 HB Nijmegen, The Netherlands. Fax: +31 24 36 10989.
E-mail address: c.beckmann@donders.ru.nl.

1053-8119/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2012.02.020

which are a mixture of true underlying source signals. Usually neither the original signals nor the mixing transformation is known — undoing this mixing process is a challenging problem known in the area of signal processing as the *blind source separation* (BSS) problem (Nandi, 1999).

Within the last 20+ years, Independent Component Analysis (ICA) has received attention from researchers in such disciplines as statistics, exploratory data analysis, signal processing and neural networks. Within the classical signal processing field ICA has been invented and reinvented over the course of decades, e.g. by looking at ICA as an extension of Principal Component Analysis (Jutten and Herault, 1991), investigating solutions to the BSS problem (Cardoso, 1989; Cardoso and Comon, 1996; Nandi, 1999) or by looking at unsupervised learning rules for solving the BSS problem based on information theoretic principles (Linsker, 1988, 1990), drawing on much earlier work on the principle of redundancy reduction (Attneave, 1954; Barlow, 1961) as a coding strategy for neurons of the perceptual system. The goal of ICA is to express a set of random variables as *linear* combinations of *statistically independent* component variables. In the context of BSS, ICA attempts to discover hidden, underlying and statistically independent source signals only from the measured observations that are *unknown* linear mixtures of *unobserved* sources (Comon, 1994).

Within the basic ICA model, we do not assume that these source distributions are known; if they are, the problem of identifying the hidden sources and the mixing is considerably simplified. In the general case of unknown source distributions both the sources and the mixing are identifiable, and thus recoverable, if and only if there exists at most one Gaussian signal among the sources (Comon, 1994). In order to achieve this decomposition, higher-order statistical moments[1] are needed. These can either be estimated explicitly as part of the unmixing procedure, or – more commonly – non-linear functions can be used to access this higher-order information. Two particularly popular approaches for ICA[2] are the Infomax algorithm (Baram and Roth, 1995; Bell and Sejnowski, 1995) and FastICA (Hyvärinen and Oja, 1997) and both approaches are based on the generic principle of using non-linear transforms of the data to drive the estimation. While the former is based on the principle of maximum information transfer,[3] the second algorithm is aimed at achieving maximum degree of non-Gaussianity for all estimated source signals. While there now exists a variety of algorithms and principled extensions that include work on non-linear, non-instantaneous (time-delayed) mixing or the incorporation of source structure (see (Roberts and Everson, 2001) or (Hyvärinen et al., 2001) for more details on the theory of ICA), these two algorithms still form the basis for many practical implementations of ICA.

(Spatial) ICA for FMRI

(McKeown et al., 1998) introduced ICA to the FMRI[4] community and proposed using a decomposition into spatially independent components in order to distinguish between non-task-related signal components, movements and other artefacts, as well as task-related activation. By looking for spatial independence, the decomposition conforms to the localisation paradigm of classical neuroscience. Originally derived from clinical experience, this paradigm is based on the observation that psycho-motor functions are performed in localised areas in the brain that can be inferred from specific deficits in patients. This naturally leads to the assumption that brain areas that respond to the psycho-motor task are independently distributed from brain areas affected by other sources of variability. It is important to note that this does not require these areas to be completely non-overlapping but only that other sources of signal change are not distributed the same way as the task-related areas, i.e. that knowledge about the spatial distribution of one does not provide any information on the spatial distribution of the other.

Fig. 1 illustrates how the data is represented in order to apply the ICA decomposition. The entire 4-dimensional data set is rearranged into a 2-dimensional matrix by arranging all voxels for each timepoint into a single row (i.e., one row per 3D functional image). This data set is then decomposed into two new matrices, the first one containing a time course of an underlying signal in each column and the second matrix containing a spatial component's map in each row. These, for instance, might be maps of stimulus-induced activity, task-unrelated ('ongoing') activity or maps of signal artefacts. The associated time courses then describe how each one of these multiple underlying effects contributes to the measured data at each measured point in time (i.e. in each brain image acquired in the functional run). The time courses are called the *source directions* or *signal signatures* of the data (Nandi, 1999) and jointly span the space of all temporal signals identified by the ICA decomposition. Thus, spatial independent component analysis can be viewed as a way of finding *temporal* basis vectors so that the associated spatial maps are sparse and statistically independent. The similarity with the General Linear Model (GLM) is quite obvious, with the time-course matrix taking on the role of the GLM design matrix. The only fundamental difference is that instead of having to specify a design matrix prior to the analysis and then estimating the (spatial maps of) effect size parameters in the GLM, in ICA both the mixing matrix and the maps of effect sizes are being estimated simultaneously from the data, using information theoretic principles to drive the joint estimation of these two quantities.

The basic idea of splitting the data into modes on the basis of spatial independence and sparseness immediately generated debate in the field, e.g. Friston (1998) argued that even though different brain functions might be spatially localised, the principle of functional integration might imply that neuronal processes share a large proportion of cortical anatomy, rendering such a decomposition approach problematic.

Despite the ongoing discussions, the 1998 paper by McKeown and colleagues managed to significantly (re)vitalise the research area of exploratory FMRI data analysis. Various groups and individuals started evaluating spatial vs. temporal ICA (Calhoun et al., 2001b; Stone et al., 1999), different methods for estimation (Esposito et al., 2002) and extensions e.g. to constrain estimation to cortical surfaces (Formisano et al., 2004) or to incorporate paradigm information (Lin et al., 2010).

Further, and in parallel with developments in GLM modelling, the field has seen a variety of different approaches being introduced for multi-subject/multi-group ICA (Beckmann and Smith, 2005; Calhoun et al., 2001a, 2008; Esposito et al., 2005; Guo and Pagnoni, 2008; Svensén et al., 2002). With the release of dedicated software tools (Brainvoyager (2000), FSL (2001), DTU Toolbox (2002), GIFT (2004)), ICA started to become available to the wider non-methods community of clinical and cognitive neuroscientists, leading to a steady increase in the number of publications using ICA for part of the image analysis (see Fig. 2).

Probabilistic ICA and MELODIC

My personal involvement in the area of ICA/FMRI research started in early 1999 when I was fortunate to join the FMRIB Centre in Oxford to start working towards a DPhil in Information Engineering. I came

[1] i.e. statistical quantities other than the mean and variance, such as skew and kurtosis.

[2] In general, and ICA for FMRI specifically.

[3] Or, equivalently, minimization of mutual information between estimated sources.

[4] In addition to the reasons listed in Jenkinson et al. (2012-this issue) I am particularly determined to use the upper-case F. The tools and techniques employed in the statistical analysis of functional data are closely related to classical time-series analysis and are quite different from standard image processing/ computer vision techniques that are the bread-and-butter of structural MR analysis.

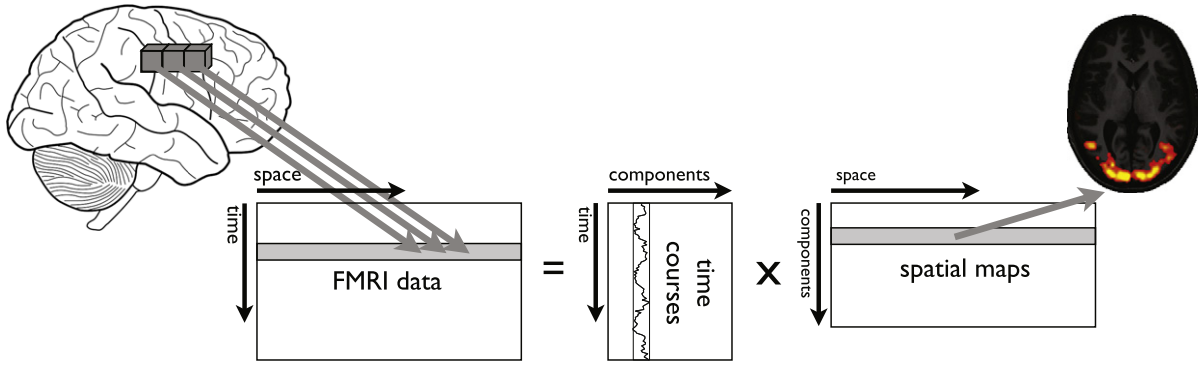

Fig. 1. Schematic illustration of the data representation and the spatial decomposition performed by spatial ICA on FMRI data.

with a background in Maths and (a bit of) applied Statistics from various places and had spent some time in industry, modelling financial time series data. I had first come across ICA in 1996 when working towards an MSc in the Math and Stat departments in Oxford and was intrigued by its effectiveness in finding hidden signals. The FMRIB centre itself opened in 1998 and things were still in the process of starting up. Steve Smith was Head of Analysis — the powers in Oxford at the time fortunately did realise that much innovation is to come from dedicated physics and analysis research teams with their own research programmes. I joined a small team of people there – the early beginnings of the general analysis efforts are reported elsewhere in this issue (Jenkinson et al., 2012-this issue) – on a DPhil project on *Independent Component Analysis for FMRI*.

During these early days at FMRIB there was an overarching theme of focussing on probabilistic models and improved statistics both for functional and structural analysis, ultimately resulting in e.g. developing a probabilistic tractography approach (Behrens et al., 2003) and improving on GLM modelling by incorporating time-series pre-whitening (Woolrich et al., 2001). The ICA work, therefore, was similarly aiming at placing the model into a probabilistic framework. In the first instance this amounted to augmenting the classical noise-free ICA model with an explicit stochastic noise term. A variety of the approaches tried and combined in the PICA framework closely relate to ideas a number of us 'boys' were looking at during these early days. For instance, the voxel-wise pre-whitening approach that Mark Woolrich worked on ended up being investigated as part of the initial data pre-processing loop in PICA (being calculated from the initial noise estimates following the probabilistic PCA step).

Fig. 3 shows[5] the flow diagram of the probabilistic ICA model as presented in (Beckmann and Smith, 2004). The original data is first demeaned and normalised to have unit *noise* variance, a process we termed *variance normalisation* (VN). This step is designed to ensure that in the absence of any non-Gaussian signals the ICA decomposition does not simply end up getting drawn towards voxels in the brain that show uninteresting (Gaussian distributed) strong signals (such as voxels within the cerebrospinal fluid which typically are very bright and have lots of associated variance). If this is not controlled for, ICA will not – in the absence of signal – have a uniform false-positive detection rate everywhere. This non-homogeneity across space had been observed a while earlier (see McKeown and Sejnowski, 1998, Fig. 1), though in the context of describing the likelihood of detecting activation. There is a bit of a chicken and egg problem here, in that this variance normalisation requires knowledge of the signal in order to define what the noise is that needs to be forced to have the same noise variance across space. The solution proposed in Beckmann and Smith (2004) is to iteratively perform normalisation and probabilistic PCA (PPCA) to split the total data space into initial noise and signal sub-spaces where the former can then be used to reiterate and refine the normalisation steps. The independent component maps can then be estimated from the pre-processed data. In the case of MELODIC, the unmixing is performed using the FASTICA technique, i.e. using a fixed-point iteration scheme to optimise for non-Gaussianity by maximising the neg-entropy of the signals (Hyvärinen, 1997).

This model contains some other innovations, e.g. the introduction of the idea that after having removed the inherent spatial bias towards high-variance voxels by means of VN one can introduce explicit spatial bias quite simply by means of modifying the calculation of the data covariance matrix used by the PPCA step. For instance, the calculation of this covariance can be restricted to only include voxel pairs within a certain (Euclidean or cortical surface) distance of each other or restrict/weight the calculation by the probability of voxel pairs to belonging to the same tissue type. The options are plentiful and much of these ideas remain under-explored to this day.

The final step of the PICA procedure is concerned with the process of statistical inference. The estimated IC maps are transformed to voxel-wise *Z*-statistics by dividing the raw IC maps estimated by ICA by the standard deviation of the residuals from the initial PPCA.

Finally, these maps are thresholded in order to infer voxels that are significantly modulated by the component time courses. This is done using a mixture model fitted to the component's histogram. The ideas that led to this thresholding approach were discussed within our group to be used also for traditional (GLM-based) analysis — both Mark Woolrich and I ended up being intrigued by the possibilities offered by mixture modelling of effect size maps or statistical images as proposed by Everitt and Bullmore (1999) and in particular the Gaussian/Gamma mixture approach proposed by

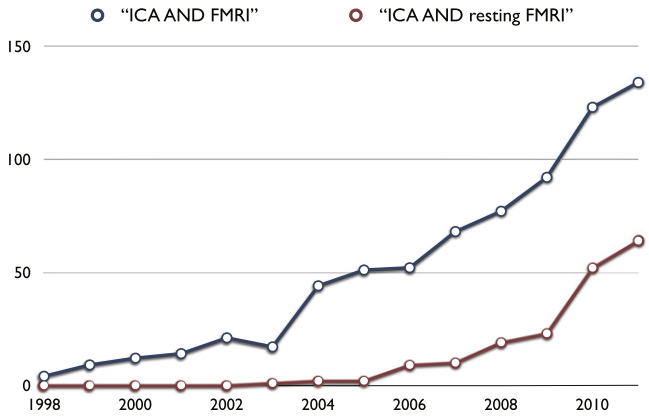

Fig. 2. Annual number of publications that contain the keywords "ICA AND FMRI" or "ICA AND resting FMRI" in their title, abstract or keywords (search on www.scopus.com, accessed 28.11.2011).

[5] I clearly remember a discussion with Steve on the benefit of having such a flow diagram relative to the amount of time it took to generate. I now draw great pleasure out of re-using this one figure (almost unaltered) for the fourth time....

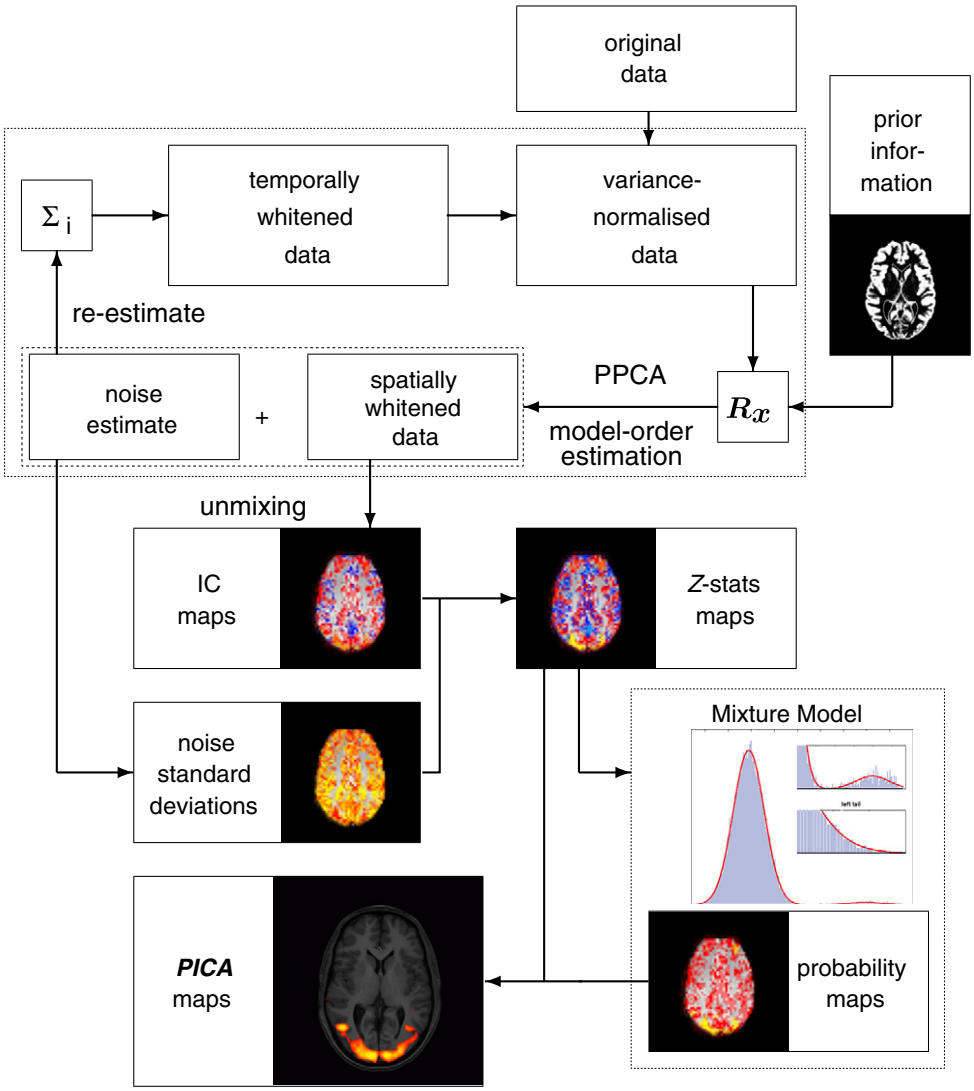

Fig. 3. Schematic illustration of the probabilistic ICA framework (Beckmann and Smith, 2004).

Hartvig and Jensen (2000). Mark at the time went the extra mile of developing a spatial mixture model with adaptive regularisation to be used on FMRI activation/statistic maps, where neighbourhood information is explicitly incorporated into the estimation of class membership (Woolrich et al., 2005). In MELODIC I simply made do with a non-spatial version of the Gaussian/Gamma mixture model. This was not just due to pure laziness on my part. The main reason was that I perceived ICA to be a powerful spatial filter already and therefore felt that any further spatial regularisation (such as one induced by having a Markov Field regularising the voxel-wise class memberships) would no longer be necessary and – in fact – be detrimental in cases where the spatial maps show strong edges. After ICA unmixing, the source histograms typically all show very clear Gaussian background noise densities, aiding the mixture model fitting. A second reason for sticking with the non-spatial variant was the vast amount of variation in spatial characteristics across the artefact maps estimated by ICA (see also Fig. 5) — spatial regularisation would inevitably smooth across some of these artefacts, making their source density look more like those from BOLD signals, making a clear separation of signals of interest from artefacts more difficult. Once fitted, this approach permits thresholding beyond simple null-hypothesis testing and employs an alternative hypothesis testing approach: because both distributions for the background noise and the signal get modelled explicitly, one can threshold based on the relative chance of a voxel's intensity being more probable under the 'null' or under the 'alternative'.

The quality of the fit, and therefore the utility of this approach for thresholding ICA maps, interacts strongly with the initial data preprocessing. The normalisation of voxel-wise residual noise (VN) in particular is important for ensuring that the voxels which do not contain any signal end up being modelled well by a single Gaussian distribution. Fig. 4 shows an example using the simulated data from Beckmann and Smith (2004). Without variance normalisation (top) the estimated spatial map histogram is more sharply peaked than a simple Gaussian bell curve — the different noise variances associated largely with the three tissue types render this background noise distribution closer to a superposition of three Gaussians. As a result, in a model of only a single Gaussian for the background noise, the Gamma distributions – designed to model activations and deactivations in the tail of the distribution – also end up also fitting to the background noise bulk of the histogram (Fig. 4a). With appropriate variance normalisation, where each voxel's time course gets normalised to the same background variance independent of tissue type, the model fit is much improved. This, in turn, significantly reduces the false-positive rate. Essentially, after normalisation the ICA unmixing operates on initial *Z*-statistics (data normalised by noise standard deviation), so due to this pre-processing the estimation now becomes an exercise in classification (rather than regression). As a result, the

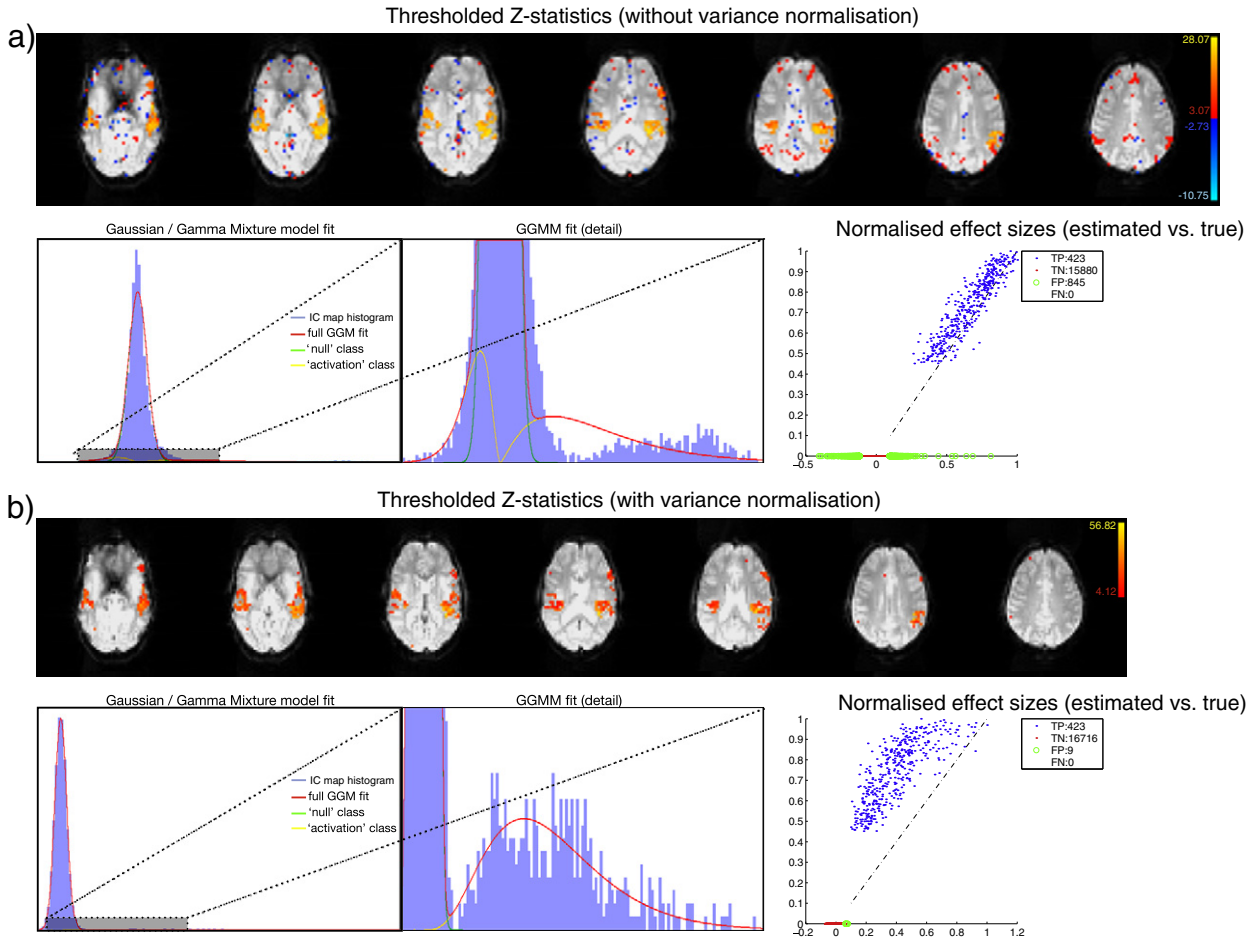

Fig. 4. Illustration of the effects of variance-normalisation and the interaction with mixture-model based thresholding: The two panels show results for non-variance normalised data (a, top) and for variance-normalised data (b, bottom) of *Z*-statistic images and their histograms from IC maps estimated from simulated data (see (Beckmann and Smith, 2004) for details). In the case of no normalisation, the 'background' noise histogram is much more sharply peaked than a simple Gaussian distribution. The overall mixture model fit can be quite poor and the Gamma distribution, intended to model the active class, also model the fat tails of the background noise class (see GGMM plots). Shown also are the scatterplots of normalised estimated maps vs. 'true' activation magnitude, together with the number of true/false positives/negatives (TP/TN/FP/FN).

scatterplot of estimated vs real effect size (Fig. 4, right) should no longer be close to a straight line fit $y = x$ but instead ideally start approximating a step function: at some SNR level the classification into active vs non-active becomes very certain. Again, this is analogous to the case of the GLM where statistical maps are not to be understood as effect size estimates. If so desired, accurate effect size estimates can be generated trivially by means of regression, once accurate classification has been achieved (i.e. masking the non *Z*-transformed IC maps by the thresholded and binarised *Z*-statistic image).

The probabilistic ICA model got implemented and released as part of the FMRIB Software Library (FSL (Smith et al., 2004; Woolrich et al., 2009), see Jenkinson et al., 2012-this issue for some history on FSL) in a basic form in June 2001 as MELODIC — Multivariate Exploratory Linear Optimised Decomposition into Independent Components. Mixture-model based inference was added in FSL3.0 (2002) and further extensions for group analysis (concat-ICA and Tensor-ICA (Beckmann and Smith, 2005)) were added as MELODIC 3.0 in FSL4.0, released in 2008.

Loose methods ends

Statistical inference

The alternative hypothesis testing approach promoted above provides a stringent framework for thresholding ICA maps. Nevertheless, it does create confusion, largely because the default output (of reporting maps $p > 0.5$) and the associated ideas of testing an alternative hypothesis explicitly against a modelled null hypothesis remain to be underrepresented in classical stats training. The field of FMRI statistics for many years has been dominated by the mantra of controlling for *false positives*, a reflection of the dominance of cognitive neuroscience in the field. It is important to note, however, that with an increasing uptake of functional imaging in the clinical domain a shift towards controlling for *false negative detections* gains equal relevance (see e.g. Bartsch et al., 2006 for examples). Ultimately, thresholding techniques should allow for explicit control of the personal loss function, i.e. a specific specification of the personal preference for false positive relative to false negative decisions. While this is a fundamental feature of the mixture model approach, the correctness of the thresholding decision now depends on the appropriateness of both the null and alternative distributions. In the case of a Gaussian/Gamma mixture model, the model for the background noise can easily be justified on the basis of generic assumptions about the FMRI noise. The Gamma model for the 'activation' classes, however, is only poorly motivated and chosen largely for mathematical convenience. One possibility is to use non-parametric characterisations of these classes, e.g. splines such as in Efron (2004). An alternative might be to instead resort to (pseudo-) null hypothesis testing, e.g. by using the Gaussian/Gamma model only to obtain a tight fit of the background noise null and resorting to null-hypothesis testing for thresholding purposes, calculating the expected false discovery rate relative to this (now well fitting) Gaussian 'null' distribution. Yet

another alternative might be to move away from thresholding altogether, maybe only enhancing clusters that show up in component maps, e.g. using techniques like *threshold-free cluster enhancement* (Smith and Nichols, 2009).

Model order selection and the intrinsic dimensionality

If questions by reviewers and on the FSL email support list are anything to go by, then the question of the number of components to be extracted from the data remains a prominent one.

While early ICA research simply extracted as many components as required to model a certain proportion of variability in the data (e.g. choose the number so as to retain 99.9% of the variability in the data), the field quickly moved towards using information theoretic estimates of the *intrinsic dimensionality* of the data (Beckmann and Smith, 2004; Beckmann et al., 2001; Calhoun et al., 2001b; Li et al., 2007; Nandi, 1999). In MELODIC, the proposed approach is to base the number of estimated components on some information theoretic analysis of the Eigen spectrum of the data covariance matrix. The question as to how to estimate the (or an) optimal source number remains actively debated. While it is easy to evaluate the quality of different approaches with simulated data, for real data this is much less trivial. It is certainly possible to define estimators that are insensitive to e.g. temporal autocorrelation of the data (Nandi, 1999), assuming that it is valid to ignore the temporal autocorrelation of the signals contained in the data along the way. The fundamental question that remains, however, is to what extent there truly is an 'optimal' number of components and what this criterion of 'optimality' would need to look like. There certainly is an issue that relates to using a linear decomposition technique like ICA for identifying signals which might not be well represented using linear approaches.[6] However, there also is increasing evidence that ICA can provide biologically interpretable decompositions across a wide range of dimensionalities, e.g. Abou-Elseoud et al. (2010) argue that decompositions across 70 ± 10 components give detailed, yet robust decompositions and in Smith et al. (2009) we demonstrate that across various dimensionalities for the analysis, different yet plausible decompositions are obtained, at the higher dimensionality of 70 describing essentially the same systems as in the 20-dimensional decomposition, just in a more fine-grained way, e.g. robustly splitting the left and right sensory–motor system or fractionating the visual system into known subdivisions such as MT etc. The difficulty in defining simple metrics for an 'optimal' dimensionality, therefore, is likely not simply due to modelling inadequacies such as assumptions of linearity but are an actual reflection of the true biological complexity of the underlying signals and systems. As we start to understand more of the biological validity of *differences* within systems (both within and across subjects) we might gain more confidence in the way that ICA splits systems into finer sub-systems and the interpretability of high-model order decompositions.[7] Neuroscientists with a strong belief in tightly controlled experimental protocols and those opposed to data exploration might shake their heads vigorously at the possibility of generating substantially different results from the same data. It is important to keep in mind, however, that this is not fundamentally different from the (typically not discussed or explored) variability in GLM results due to a plethora of modelling decisions, e.g. whether to model interactions, confounds or behavioural information at the group level etc. In this context the set of papers published as part of the first Functional Image Analysis Contest (FIAC) organised at the Human Brain Mapping conference in 2005 might serve as an informative example of the variability seen across more traditional model based approaches, when applied to the same set of data generated under a well described experimental protocol (Aston et al., 2006; Beckmann et al., 2006; Goebel et al., 2006; Poline et al., 2006; Saad et al., 2006; Suckling et al., 2006; Taylor and Worsley, 2006).

The 'killer applications'

Given the overall simplicity of the ICA model and its close relationship to the widely used GLM there clearly are many potential applications, both in clinical and cognitive neurosciences, where the added flexibility of signal modelling in ICA can be useful. For task-based studies, examples include natural-movie-viewing paradigms (Bartels and Zeki, 2004) or simulated driving studies (see (Calhoun and Pearlson, 2012) for a recent review). Nevertheless, during the first few years very few studies[8] employed ICA as the primary analysis approach. This, in large part, is due to the fact that traditional analysis (mainly within the framework of the GLM) continues to serve the community well, allowing one to ask very specific questions about the locality of significant signal changes in response to well-designed experimental manipulations.

There are, however, (at least) two 'killer apps' — application domains where ICA has been shown to be an immensely powerful tool, providing utility beyond what can be achieved with conventional methods.

Artefacts and data denoising

The first application is that of artefact detection, artefact characterisation and possibly even data denoising. FMRI is a tremendously challenging technique where complex underlying physiology is measured by highly complicated technical processes. FMRI data preprocessing aims at removing the worst of such effects, using various data processing techniques such as motion and distortion correction, temporal and spatial filtering prior to the analysis or explicit outlier modelling as part of the statistical procedure of effect size estimation. Nevertheless, empirical evidence suggests that residual noise has both spatial and temporal structures. Fig. 5 shows different kinds of artefacts that consistently appear in FMRI data (shown as a few example axial slices together with the associated example time courses (Beckmann et al., 2000)).

Artefacts arising from problems with the scanning hardware (a, d) are – because of the intricate nature of the imaging process – hard to identify and understand fully. Many technical components are involved in the image generation process and consequently there are many different possible stages at which even slight deviation from expected performance of the individual components can result in large differences in the measurements. The top image shows three instances of what is often called a slice dropout and which is caused by small gradient waveform corruptions, resulting in a ghosting of objects diagonally in the image. The artefacts at the bottom are caused by an instability in the relative timing of the slice select gradient and the radio-frequency pulse waveform, in this case due to a faulty device on the driving console. In both cases, there was a mean intensity change over the brain volume of approximately 5–10%, exceeding (and possibly swamping) any real activation of interest.

Fig. 5(b) shows spatial components that correspond to head motion of the subject. The motion of the entire object during acquisition generally results in a blurring of the entire image with ghost images in the phase encoding direction. A solution to help reduce motion artefacts is to immobilise the subjects head. However, motion of the

[6] Head motion springs to mind, where non-linear motion effects will get distributed across multiple maps — in an analogous way that non-linear functions get represented by a superposition of locally linear functions in a Taylor decomposition.

[7] This point, interestingly, relates to an issue already discussed in the original 1998 ICA for FMRI paper by McKeown and colleagues, where much space was spent discussing 'transient task-related components' relative to 'consistently task related components'.

[8] Studies other than those conducted by research groups also being heavily involved in the development of that method itself....

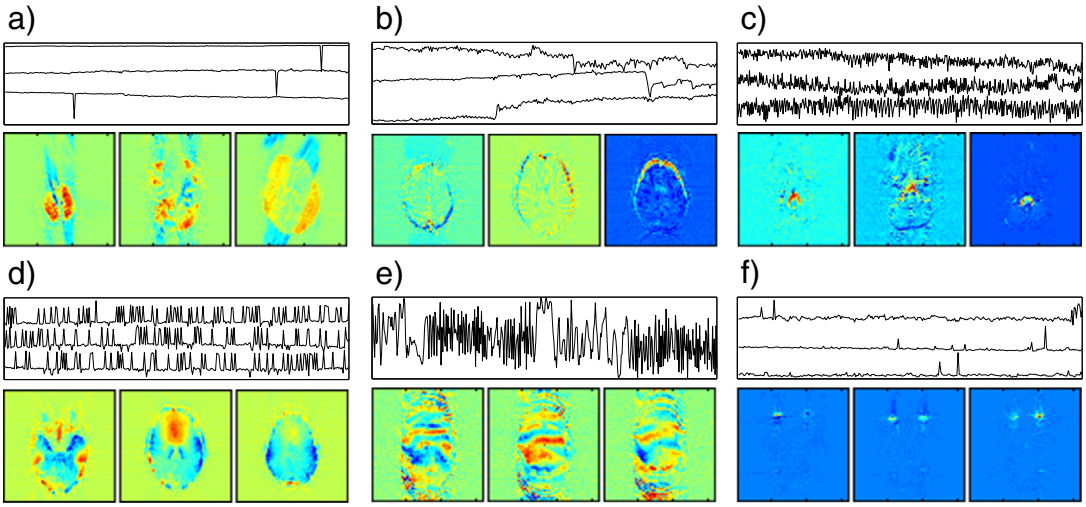

Fig. 5. A selection of 'typical' image artefacts in the eyes of ICA. These are 'typical' only in the sense of having seen these effects repeatedly in data I have been supplied with by colleagues and collaborators. It is quite likely that data with very different characteristics (e.g. different acquisition types – spirals maybe – or data acquired simultaneously along with EEG data) shows substantially different types of artefacts. Each figure shows 3 example axial slices of a spatial ICA component map, together with a selection of 'typical' associated time courses: (a) slice dropouts, (b) head motion, (c) high-frequency noise, largely within CSF, (d) hardware spiking artefact, (e) *N*/2-ghost image, interacting with head motion and (f) eye-blink artefacts.

brain can also be caused by respiration or the cardiac cycle, neither of which can easily be eliminated. In ICA component maps such artefacts typically appear as a rim-like effect near the strong intensity edges of the image (e.g. the brain/non-brain boundary) in the direction of dominant motion; in these examples the associated time courses clearly exhibit abrupt level changes, indicating a quick positional change rather than a slow gradual movement. The bottom figure shows a particularly severe case of *N*/2 ghosting.[9] This type of artefact is specific to EPI and is caused by inconsistencies in timing between the image acquisition and the gradient switching resulting in differences between even and odd acquisition lines in *k*-space. The spatial maps clearly pick out signal outside of the head (image ghost) and interference patterns show up where the ghost images overlap the head volume. The associated time course is of high frequency with additional confounds caused by motion.

A different type of high frequency noise is shown in (c), where the spatial maps identify high intensity rapid variation within the CSF in the ventricular system. The bottom figure (f) shows spatial maps and associated time courses of eye-related artefacts caused by either eye motion or eye-blink. The eyes appear clearly in the spatial maps together with vertical patterns that show ghosting along the phase-encode direction in the image. Eye blink itself is not a serious artefact for FMRI as it is possible to eliminate the signals associated with the volume of the eyes, using appropriate spatial pre-processing such as brain extraction. In some cases, however, eye-blink can combine with the N/2 ghost such that the amplitude modulation due to eye-blink wraps around the field of view into posterior cortical areas. If eye-blink is correlated with the external stimulus, this can induce additional false positives within these cortical areas in a GLM analysis.

Automatically identifying and classifying such artefacts remain a tremendous challenge. Tohka et al. (2008) use a classifier defined by a decision tree, reducing mis-classification rate (relative to a human expert) down to somewhere between 20% and 30%. This, in many cases, remains too high to be useful in practice. The reason why it remains difficult to automate classification relates to the wide variety of sources that induce such image artefacts and it is quite likely that further improvements will require much larger and better sets of hand-classified ICA decomposition across different MR sequences, field strength etc. in order to train sensitive and specific classifiers. One early attempt to generate such a database of 'typical' artefacts we initiated early in 2002 as the 'Little FMRI Shop of Horror', a web-based repository of hand-labelled ICA components. However, despite active lobbying during our FSL courses and on the mailing list we have failed to generate a substantial amount of external submissions. More work is currently underway as part of the *Human Connectome Project*, attempting to come closer to fully automated approaches for ICA-based FMRI artefact classification.

Data denoising

For statistical inference in the case of GLM analysis to be valid, the residual errors should be Gaussian distributed any non-Gaussian process that is not modelled within the design or removed during pre-processing violates this assumption. If the artefactual time course is orthogonal to the signals of interest, the presence of the artefact will not affect the GLM regression coefficients, but will increase the residual variance estimate and consequently will result in reduced statistical significance for any effect of interest (inflated false-negative rate). If, however, the artefactual time course is partially correlated with the design, it will also lead to incorrect regression coefficients. In the case of positive correlation, this will inflate the false-positive rate of the analysis. For this reason, the real 'killer application' for ICA is to use the information represented in the decomposition to reduce the negative effects of artefacts for standard GLM-based analysis. There are various possible approaches. Firstly, information from the spatial maps can be used to remove certain voxels from further analysis. Secondly, it is possible to utilise the time course information in order to identify certain points in time that require attention. In the case of slice dropouts, the affected scans can either be excluded from further analysis or the intensity value at a time point can be adjusted to the mean intensity of the volumes acquired before and after the dropout occurred. Similarly, one can think of using both spatial maps and time courses, e.g. to verify the effectiveness of standard motion correction methods by comparing ICA analysis before and after correction (Bannister et al., 2002). Finally, the most promising approach is to use ICA as an 'intelligent filter' by combining spatial maps with their associated time-courses to form an estimate of the non-Gaussian structured noise in the data, which can then simply be subtracted from the data. For one structured noise component only, this effectively results in subtracting a rescaled version of the time course of the structured noise component from each voxel's time course. The value of the voxel within the

[9] This is also often referred to as *Nyquist ghost*, even though it is not actually a problem related to the Nyquist sampling rate.

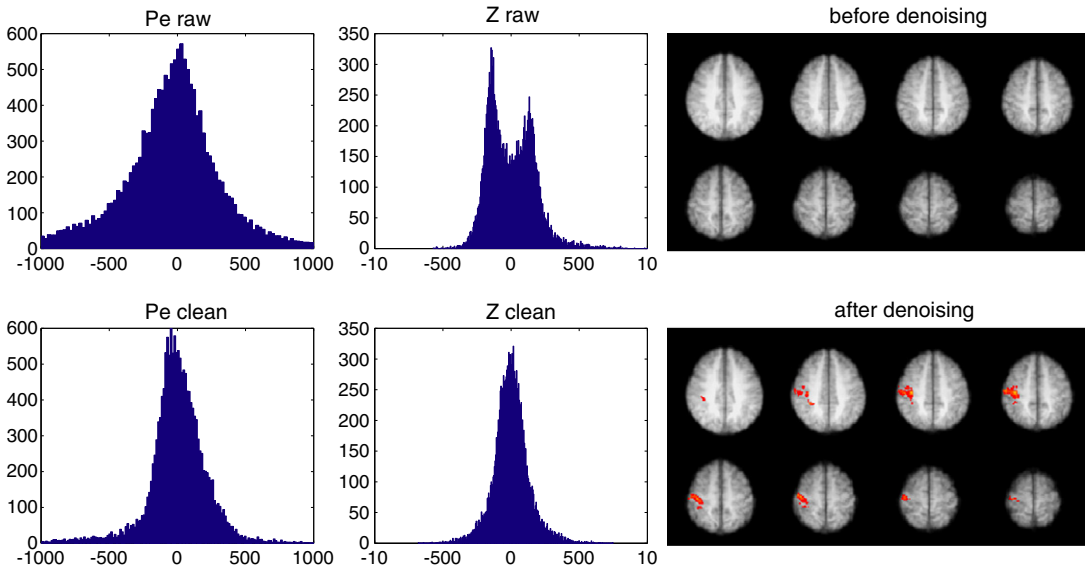

Fig. 6. Illustration of ICA-based data denoising: the top row shows the histogram of parameter estimates and Z-statistics for the 'left hand finger tapping' contrast from a bilateral manual finger tapping task. Particularly the latter points to a heavily confounded data set — the histogram looks nowhere near sensible, assuming signal sparseness. Consequently, the 'left–right' finger tapping contrast failed to identify a sensible activation pattern. Upon removal of artefacts identified by ICA, both the effect size/*Z*-statistic histograms and the ultimate thresholded contrast map are much improved.

component map determines the amplitude of the structured noise sequence being subtracted. An alternative to subtracting the estimated structured noise is to reconstruct the data from the components *not* classified as structured noise (Perlbarg et al., 2007). Note, however, that this requires the FMRI signals of interest to be contained within these non-noise estimates — that is, this approach is at odds with later testing effects under the 'null-hypothesis' as this initial denoising stage implicitly requires the alternative hypothesis of there being signal in the data to be true. Removing structured noise, on the other hand makes no assumptions about the signals of interest and therefore can be combined with later 'null hypothesis' testing, e.g. by feeding the denoised FMRI data into a classical GLM analysis.

An example of this approach is shown in Fig. 6. The data is from a subject performing left and right-sided finger movements. The original GLM analysis (top) did not show any sensory–motor activity in the 'left–right' contrast image, despite such a paradigm normally being very robust, even at the level of an individual subject. The unthresholded *Z*-statistic image for the left-finger tapping contrast (top centre) showed a disturbing degree of deviation from a simple Gaussian shape. Under the assumption of sparseness of the signal, the raw *Z*-statistic histogram should be dominated by Gaussianised background noise, i.e. should be dominated by a simple Gaussian bell curve. Upon manually identifying artefact components and removal of those, both the histogram of effect size estimates (bottom left) and the *Z*-stat histogram (bottom centre) showed much improved characteristics. The pattern of activation changed significantly. The 'left–right' contrast image shows clear and biologically plausible 'activation' within the right sensory–motor system.

Resting-state FMRI

The second 'killer application' of ICA is that of estimating patterns of functional connectivity from FMRI data acquired under a resting condition (Biswal et al., 1995; Fox and Raichle, 2007; Raichle et al., 2001). Data acquired with resting-FMRI does not lend itself easily to a standard GLM analysis. Seed-based correlation analysis, first suggested by Biswal et al. (1995), and ICA have therefore emerged as the two dominant analysis approaches. For ICA and FMRI, an increasing proportion of work is research into resting state functional connectivity, now accounting for roughly 50% of the overall published research (Fig. 2).

The first application of ICA to resting FMRI data dates back to 1998/99 (Cordes et al., 1999) though the first published paper did not appear until 2003 (Kiviniemi et al., 2003) — which likely reflects a certain level of scepticism on the part of reviewers. Indeed, the overall utility of studies into the resting state has been debated for a substantial amount of time, largely considering the utility from a cognitive experimental neuroscience perspective (Morcom and Fletcher, 2007a,b; Raichle and Snyder, 2007).

The FMRIB foray into this area initially was motivated by identifying the presence of these effects across a variety of task-activation data. The immediate concern was again very much focused on the statistical implications for task FMRI analysis. In the late 90s there continued to be a debate about 'global intensity normalisation', the process of scaling every volume of the 4D sequence in order to achieve a (temporally) constant mean spatial intensity. Such processing was common (and beneficial) for the analysis of PET data and many of these approaches ended up being recommended for FMRI data analysis, too (Friston et al., 1994). Along with the pre-whitening vs. pre-colouring debate (see (Woolrich et al., 2001)), the question of the global intensity normalisation generated significant amount of attention (Della-Maggiore et al., 2002). One of FMRIB's first contributions to the field was the realisation that the issue of the global mean signal and the presence of Resting-State Networks even in task data is heavily intertwined. Fig. 7 shows an example figure (modified from DeLuca et al., 2002a) demonstrating that the global mean signal across all voxels, the global mean across grey-matter only and specific resting-state network time courses share a common power spectral density, demonstrating that the 'global' mean signal (calculating by simply averaging across all voxels) actually has an associated non-global spatial structure (estimable by regressing the data back onto the global mean time course) which reflects the presence of multiple resting-state networks. These results indicate that in fact much of what is seen as physiological noise in task FMRI can be attributed to resting-state networks. Interestingly,[10] the arguments pro and against 'global intensity normalisation' as a pre-processing

[10] ...even recent history seems to repeat itself....

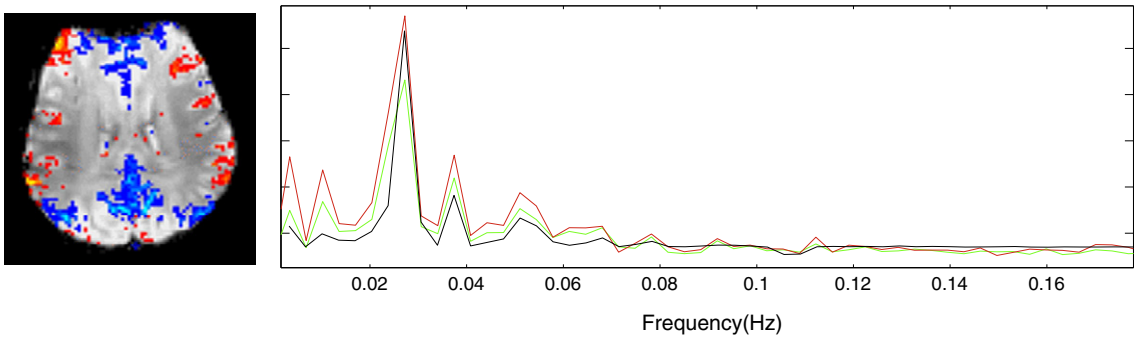

Fig. 7. Resting-state network with detail (left) identified using PICA. The plot shows the power spectra from a PICA derived RSN time course (black), the mean intensity time course (MIT) for grey-matter voxels (red) and the MIT for the entire brain (green) (see (DeLuca et al., 2002a) for details).

step have re-appeared in recent years in the context of 'Global Signal Regression' for seed-based analysis of resting FMRI data (Cole et al., 2010; Fox and Raichle, 2007; Murphy et al., 2009). Our research quickly started to appreciate the importance of RSNs over and above their role as 'noise' — attempting to explore the utility of ICA for RSN identification (Beckmann et al., 2005; DeLuca et al., 2002b), characterising the robustness of these effects (Damoiseaux et al., 2006) and, more recently, their role in cognitive (e.g. (Smith et al., 2009)) and clinical (e.g. (Filippini et al., 2009)) neurosciences.

These investigations, along with other work on RSNs and ICA have greatly increased knowledge about resting functional connectivity, e.g. by establishing high test–retest reliability of ICA-based identification of RSNs (Zuo et al., 2010), demonstrating clinical utility as sensitive markers of neurodegenerative disease (Greicius et al., 2004), adding biological validity to RSNs by means of demonstrating significant heritability (Glahn et al., 2010) or demonstrating that ICA provides sensitive biomarkers for psychiatric illnesses characterised not only by disrupted local brain connectivity within a single RSN but also in terms of disrupted global connectivity between resting-state systems (Calhoun et al., 2009). Relative to the dominant alternative (seed-based) view on functional connectivity it is the inherent multivariate nature of ICA that provides the means to *simultaneously* detect and characterise various different RSNs from a short, single session of resting data. This helps to shape a view on the brain's functional connectome. Both ICA and resting FMRI will likely form a crucial part of the *Human Connectome Project* (Sporns, 2011), aimed at mapping the complex hierarchical organisation of functional and structural pathways that underlie brain function.

Future directions

It is impossible to know what the next big thing is going to be for the use of ICA in FMRI. I will take the liberty of speculating, based on my own personal views on what I currently see as interesting possible avenues for further development.

Variability and individual differences

The inherent un(der)constrained nature of ICA offers a great opportunity to look specifically at the between-subject variability in greater detail. In many cases, the observed variability in ICA estimates is interpreted negatively in the sense of this being a reflection of a modelling inadequacy. In many application domains, however, such estimated cross-subject differences might actually be important in that they reflect true (and hopefully ultimately interpretable) biologically meaningful differences, e.g. might signify development of or decline within functional systems, signify differences in genetic makeup or pick up on differences in functional strategy employed for the performance of a given task. Understanding these differences, therefore, can be extremely valuable for furthering the utility of imaging neurosciences. Developing methodology that can accurately characterise those differences at a population level, i.e. model common features along with non-common specific variations thereof, remains a tremendous challenge. It will require the development of new tools and techniques, but potentially will provide powerful ways of looking for data driven population stratification, e.g. where on the basis of individual differences in one or multiple estimated signals a patient population that might appear homogeneous on the basis of behaviour and clinical presentation gets characterised into sub-populations that might correspond to real biological differences in the underlying disease mechanisms.

Multi-modal imaging

Typical imaging sessions acquire a wide range of structural and functional information non-invasively, thereby generating a vast amount of data. In addition, a plethora of auxiliary information (behavioural, physiological, questionnaire data, genetic data) often is available along with the imaging data. More work is needed to handle the integration of information across these different types of data. One approach, attractive for its conceptual simplicity, is to simply put all this data into a single matrix and feed this hybrid matrix into a single ICA decomposition (Calhoun et al., 2006). Within this approach it is difficult to appropriately weight the potentially vastly different degrees of 'informativeness' contained in each data modality (due to e.g. differences in the number of samples, their units, modality-specific signal-to-noise characteristics, modality-specific sparseness, etc.) so simple data concatenation is not likely to provide the answer to the question of how to make use of this data explosion. The alternative of running separate ICA decompositions similarly might not easily lend itself to post-hoc data fusion across the different sets of components. Recently Groves et al. (2011) introduced a novel 'Linked-ICA' model for simultaneously modelling common features across multiple data modalities. Clearly, more work is required to fully explore and harness the opportunities that such advanced modelling affords.

Temporal dynamics and ICA

The application of ICA in FMRI has quickly settled on *spatial* decompositions. In large part this is due to the fact that typical FMRI data has many more voxels than time points, enabling better estimates for higher-order statistics/non-Gaussianity across the spatial domain than is possible across only a few hundreds of time points. This results in decompositions where individual component maps are largely non-overlapping. The degree of *functional* interactivity between different components, however, remains under explored. Recent advances in MR sequence development (Feinberg et al., 2010)

have opened the door to short-TR functional imaging so as to overcome limitations of low number of samples in the temporal domain. This, in turn, permits more advanced investigations into the rich temporal dynamics and interactions between source signals. Identifying temporally independent modes of function amongst spatially independent signals might shed more light onto the fundamental building blocks of brain activity, potentially providing the missing link between structure and function. Again, further methodological improvements will be required to utilise this newly emerging data more fully.

Conclusion

Independent Component Analysis has shown great utility both in cognitive and clinical imaging neurosciences. There remain a variety of methodological, conceptual and practical issues that need to be addressed in order to fully utilise the power of ICA in imaging neuroscience. Nevertheless, this exploratory approach to FMRI analysis has already been shown to provide an important complementary tool, e.g. helping to characterise brain function even in the absence of experimental manipulations. As imaging neurosciences start to embrace new imaging paradigms and to add more data modalities into their investigations in order to address increasingly complex questions about human brain function in health and disease, the role of exploratory tools in general, and ICA in particular, can only increase. I am therefore hopeful that in another 20 years, more exciting research on ICA for FMRI will be reported on.

Acknowledgments

I am grateful to Steve Smith – amongst many things – for helpful comments on this manuscript.

References

Abou-Elseoud, A., Starck, T., Remes, J., Nikkinen, J., Tervonen, O., Kiviniemi, V.J., 2010. The effect of model order selection in group PICA. Hum. Brain Mapp. 31 (8), 1207–1216.
Aston, J.A.D., Turkheimer, F.E., Brett, M., 2006. HBM functional imaging analysis contest data analysis in wavelet space. Hum. Brain Mapp. 27 (5), 372–379.
Attneave, F., 1954. Some informational aspects of visual perception. Psychol. Rev. 61, 183–193.
Bannister, P.R., Beckmann, C.F., Jenkinson, M., 2002. Motion artefact decorrelation in FMRI analysis using ICA. Proc. Int. Soc. of Magnetic Resonance in Medicine.
Baram, Y., Roth, Z., 1995. Forecasting by density shaping using neural networks. Computational Intelligence for Financial Engineering, pp. 57–71.
Barlow, H., 1961. Possible Principles Underlying the Transformations of Sensory Messages. MIT press, pp. 217–234. Ch. 2.
Bartels, A., Zeki, S., 2004. The chronoarchitecture of the human brain—natural viewing conditions reveal a time-based anatomy of the brain. Neuroimage 22 (1), 419–433.
Bartsch, A.J., Homola, G., Biller, A., Solymosi, L., Bendszus, M., 2006. Diagnostic functional MRI: illustrated clinical applications and decision-making. J. Magn. Reson. Imaging 23 (6), 921–932.
Beckmann, C.F., Smith, S.M., 2004. Probabilistic independent component analysis for functional magnetic resonance imaging. IEEE Trans. Med. Imaging 23 (2), 137–152.
Beckmann, C.F., Smith, S.M., 2005. Tensorial extensions of independent component analysis for multisubject FMRI analysis. Neuroimage 25 (1), 294–311.
Beckmann, C.F., Noble, J., Smith, S.M., 2000. Artefact detection in fmri data using independent component analysis. Sixth Int. Conf. on Functional Mapping of the Human Brain, p. 614.
Beckmann, C.F., Noble, J., Smith, S.M., 2001. Investigating the intrinsic dimensionality of FMRI data for ICA. Seventh Int. Conf. on Functional Mapping of the Human Brain.
Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. Philos. Trans. R. Soc. Lond. B Biol. Sci. 360 (1457), 1001–1013.
Beckmann, C.F., Jenkinson, M., Woolrich, M.W., Behrens, T.E.J., Flitney, D.E., Devlin, J.T., Smith, S.M., 2006. Applying FSL to the FIAC data: model-based and model-free analysis of voice and sentence repetition priming. Hum. Brain Mapp. 27 (5), 380–391.
Behrens, T.E.J., Woolrich, M.W., Jenkinson, M., Johansen-Berg, H., Nunes, R.G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S.M., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magn. Reson. Med. 50 (5), 1077–1088.
Bell, A., Sejnowski, T., 1995. An information maximization approach to blind separation and blind deconvolution. Neural Comput. 7 (6), 1129–1159.
Biswal, B.B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 34, 537–541.
Calhoun, V.D., Pearlson, G.D., 2012. A selective review of simulated driving studies: combining naturalistic and hybrid paradigms, analysis approaches, and future directions. Neuroimage 59 (1), 25–35.
Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001a. A method for making group inferences from functional mri data using independent component analysis. Hum. Brain Mapp. 14 (3), 140–151.
Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001b. Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. Hum. Brain Mapp. 13 (1), 43–53.
Calhoun, V.D., Adali, T., Giuliani, N., Pekar, J., 2006. Method for multimodal analysis of independent source differences in schizophrenia. Hum. Brain Mapp. 27, 47–62.
Calhoun, V.D., Liu, J., Adali, T., 2008. A review of group ICA for fmri data and ICA for joint inference of imaging, genetic, and ERP data. Neuroimage 45, S163–S172.
Calhoun, V.D., Eichele, T., Pearlson, G., 2009. Functional brain networks in schizophrenia: a review. Front. Hum. Neurosci. 3, 17.
Cardoso, J.-F., 1989. Source separation using higher order moments. Proc. ICASSP'89, pp. 2109–2112.
Cardoso, J.-F., Comon, P., 1996. Independent component analysis, a survey of some algebraic methods. Proc. ISCAS'96, 2, pp. 93–96.
Cole, D.M., Smith, S.M., Beckmann, C.F., 2010. Advances and pitfalls in the analysis and interpretation of resting-state fmri data. Front. Syst. Neurosci. 4, 8.
Comon, P., 1994. Independent component analysis — a new concept? Signal Process. 36, 287–314.
Cordes, D., Carew, J.D., Eghbalnia, H., Meyerand, M.E., Quigley, M.A., Arfanakis, K., Assadi, A., Turski, P.A., Haughton, V.M., 1999. Resting-stats functional connectivity study using Independent Component Analysis. Proc. Int. Soc. of Magnetic Resonance in Medicine, p. 1-1.
Damoiseaux, J.S., Rombouts, S.A.R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. Proc. Natl. Acad. Sci. U.S.A. 103 (37), 13848–13853.
Della-Maggiore, V., Chau, W., Peres-Neto, P., McIntosh, A., 2002. An empirical comparison of SPM preprocessing parameters to the analysis of fMRI data. Neuroimage 17 (1), 19–28.
DeLuca, M., Beckmann, C.F., Behrens, T.E.J., Clare, S., Matthews, P.M., Stefano, N.D., Woolrich, M.W., Smith, S.M., 2002a. Low Frequency Signals in FMRI — "Resting State Networks" and the "Intensity Normalisation Problem".
DeLuca, M., Beckmann, C.F., Clare, S., Behrens, T.E.J., Stefano, N.D., Matthews, P.M., Woolrich, M.W., Smith, S.M., 2002b. Further investigations into "resting state networks" and the "intensity normalisation problem". Eighth Int. Conf. on Functional Mapping of the Human Brain.
Efron, B., 2004. Large-scale simultaneous hypothesis testing: the choice of a null hypothesis. J. Am. Stat. Assoc. 99 (465), 96–104.
Esposito, F., Formisano, E., Seifritz, E., Goebel, R., Morrone, R., Tedeschi, G., Di Salle, F., 2002. Spatial independent component analysis of functional MRI time-series: to what extent do results depend on the algorithm used? Hum. Brain Mapp. 16 (3), 146–157.
Esposito, F., Scarabino, T., Hyvärinen, A., Himberg, J., Formisano, P.T., Comani, S., Tedeschi, G., Goebel, R., Seifritz, E., Salle, F.D., 2005. Independent component analysis of fmri group studies by self-organizing clustering. Neuroimage 25 (1), 193–205.
Everitt, B., Bullmore, E., 1999. Mixture model mapping of brain activation in functional magnetic resonance images. Hum. Brain Mapp. 7, 1–14.
Feinberg, D.A., Moeller, S., Smith, S.M., Auerbach, E., Ramanna, S., Gunther, M., Glasser, M.F., Miller, K.L., Ugurbil, K., Yacoub, E., 2010. Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging. PLoS One 5 (12), e15710.
Filippini, N., Macintosh, B., Hough, M., Goodwin, G., Frisoni, G.B., Smith, S.M., Matthews, P., Beckmann, C.F., Mackay, C., 2009. Distinct patterns of brain activity in young carriers of the APOE-ε4 allele. Proc. Natl. Acad. Sci. U.S.A. 106 (17), 7209–7214.
Formisano, E., Esposito, F., Salle, F.D., Goebel, R., 2004. Cortex-based independent component analysis of fMRI time series. Magn. Reson. Imaging 22 (10), 1493–1504.
Fox, M.D., Raichle, M.E., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat. Rev. Neurosci. 8 (9), 700–711.
Friston, K.J., 1998. Modes or models: a critique on independent component analysis for fMRI. Trends Cogn. Sci. (Regul Ed) 2 (10), 373–375.
Friston, K., Jezzard, P., Turner, R., 1994. Analysis of functional MRI timeseries. Hum. Brain Mapp. 1, 153–171.
Glahn, D.C., Winkler, A.M., Kochunov, P., Almasy, L., Duggirala, R., Carless, M.A., Curran, J.C., Olvera, R.L., Laird, A.R., Smith, S.M., Beckmann, C.F., Formisano, P.T., Blangero, J., 2010. Genetic control over the resting brain. Proc. Natl. Acad. Sci. U.S.A. 107 (3), 1223–1228.
Goebel, R., Esposito, F., Formisano, E., 2006. Analysis of functional image analysis contest (FIAC) data with Brainvoyager QX: from single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. Hum. Brain Mapp. 27 (5), 392–401.
Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V., 2004. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. Proc. Natl. Acad. Sci. U.S.A. 101 (13), 4637–4642.
Groves, A.R., Beckmann, C.F., Smith, S.M., Woolrich, M.W., 2011. Linked independent component analysis for multimodal data fusion. Neuroimage 54 (3), 2198–2217.
Guo, Y., Pagnoni, G., 2008. A unified framework for group independent component analysis for multi-subject fMRI data. Neuroimage 42 (3), 1078–1093.
Hartvig, N.V., Jensen, J.L., 2000. Spatial mixture modeling of fMRI data. Hum. Brain Mapp. 11 (4), 233–248.
Hyvärinen, A., 1997. A family of fixed-point algorithms for independent component analysis. Proc. IEEE Int. Conf on Acoustics: Speech and Signal Processing (ICASSP'97), pp. 3917–3920.

Hyvärinen, A., Oja, E., 1997. A fast fixed-point algorithm for independent component analysis. Neural Comput. 9 (7), 1483–1492.
Hyvärinen, A., Karhunen, J., Oja, E., 2001. Independent Component Analysis.
Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. NeuroImage 62 (2), 781–789 (this issue).
Jutten, C., Herault, J., 1991. Blind separation of sources, part i: an adaptive algorithm based on neuromimetic architecture. Signal Process. 24, 1–10.
Kiviniemi, V.J., Kantola, J.-H., Jauhiainen, J., Hyvärinen, A., Tervonen, O., 2003. Independent component analysis of nondeterministic fMRI signal sources. Neuroimage 19, 253–260.
Li, Y.-O., Adali, T., Calhoun, V.D., 2007. A feature-selective independent component analysis method for functional MRI. Int. J. Biomed. Imaging 2007, 1–13.
Lin, Q.-H., Liu, J., Zheng, Y.-R., Liang, H., Calhoun, V.D., 2010. Semiblind spatial ICA of fMRI using spatial constraints. Hum. Brain Mapp. 31 (7), 1079–1088.
Linsker, R., 1988. Self-organization in a perceptual network. Computer 21, 105–117.
Linsker, R., 1990. Self-organization in a Perceptual System: How Network Models and Information Theory may Shed Light on Neural Organization. No. 10.
McKeown, M.J., Sejnowski, T.J., 1998. Independent component analysis of fmri data: examining the assumptions. Hum. Brain Mapp. 6 (5–6), 368–372.
McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.-P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J., 1998. Analysis of fMRI data by blind separation into independent spatial components. Hum. Brain Mapp. 6 (3), 160–188.
Morcom, A., Fletcher, P., 2007a. Does the brain have a baseline? Why we should be resisting a rest. Neuroimage 37 (4), 1073–1082.
Morcom, A., Fletcher, P., 2007b. Cognitive neuroscience: the case for design rather than default. Neuroimage 37 (4), 1097–1099.
Murphy, K., Birn, R.M., Handwerker, D.A., Jones, T.B., Bandettini, P.A., 2009. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? Neuroimage 44 (3), 893–905.
Nandi, A., 1999. Blind Estimation using Higher-order Statistics.
Perlbarg, V., Bellec, P., Anton, J.-L., Pelegriniissac, M., Doyon, J., Benali, H., 2007. CORSICA: correction of structured noise in fMRI by automatic identification of ica components. Magn. Reson. Imaging 25 (1), 35–46.
Poline, J.-B., Strother, S.C., Dehaene-Lambertz, G., Egan, G.F., Lancaster, J.L., 2006. Motivation and synthesis of the FIAC experiment: reproducibility of fmri results across expert analyses. Hum. Brain Mapp. 27 (5), 351–359.
Raichle, M.E., Snyder, A.Z., 2007. A default mode of brain function: a brief history of an evolving idea. Neuroimage 37 (4), 1083–1090.
Raichle, M.E., MacLeod, A., Snyder, A.Z., Powers, W., 2001. Inaugural article: a default mode of brain function. Proc. Natl. Acad. Sci. 98 (2), 676–682.
Roberts, S., Everson, R., 2001. Independent Component Analysis: Principles and Practice. Cambridge University Press.
Saad, Z.S., Chen, G., Reynolds, R.C., Christidis, P.P., Hammett, K.R., Bellgowan, P.S.F., Cox, R.W., 2006. Functional imaging analysis contest (FIAC) analysis according to AFNI and SUMA. Hum. Brain Mapp. 27 (5), 417–424.
Smith, S.M., Nichols, T., 2009. Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. NeuroImage 44 (1), 83–98.
Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., DeLuca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., Stefano, N.D., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage 23 (Suppl. 1), S208–S219.
Smith, S.M., Fox, P., Miller, K., Glahn, D.C., 2009. Correspondence of the brain's functional architecture during activation and rest. Proc. Natl. Acad. Sci. 106 (31), 13040–13045.
Sporns, O., 2011. The human connectome: a complex network. Ann. N. Y. Acad. Sci. 1224, 109–125.
Stone, J.V., Porrill, J., Büchel, C., Friston, K.J., 1999. Spatial, temporal, and spatiotemporal independent component analysis of fMR data. Spatio-temporal Modelling and its Applications, pp. 1–4.
Suckling, J., Davis, M.H., Ooi, C., Wink, A.M., Fadili, J., Salvador, R., Welchew, D., Sendur, L., Maxim, V., Bullmore, E.T., 2006. Permutation testing of orthogonal factorial effects in a language-processing experiment using fMRI. Hum. Brain Mapp. 27 (5), 425–433.
Svensén, M., Kruggel, F., Benali, H., 2002. ICA of fMRI group study data. Neuroimage 16, 551–563.
Taylor, J.E., Worsley, K.J., 2006. Inference for magnitudes and delays of responses in the fiac data using BRAINSTAT/FMRISTAT. Hum. Brain Mapp. 27 (5), 434–441.
Toga, A.W., Mazziotta, J.C., 2002. Brain Mapping: The Methods.
Tohka, J., Foerde, K., Aron, A.R., Tom, S.M., Toga, A.W., Poldrack, R.A., 2008. Automatic independent component labeling for artifact removal in fMRI. Neuroimage 39 (3), 1227–1245.
Woolrich, M., Ripley, B., Brady, J.M., Smith, S.M., 2001. Temporal autocorrelation in univariate linear modelling of FMRI data. Neuroimage 14 (6), 1370–1386.
Woolrich, M.W., Behrens, T.E.J., Beckmann, C.F., Smith, S.M., 2005. Mixture models with adaptive spatial regularisation for segmentation with an application to FMRI data. IEEE Trans. Med. Imaging 24 (1), 1–11.
Woolrich, M.W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T.E.J., Beckmann, C.F., Jenkinson, M., Smith, S.M., 2009. Bayesian analysis of neuroimaging data in FSL. Neuroimage 45 (1 Suppl.), S173–S186.
Zuo, X.-N., Kelly, C., Adelstein, J.S., Klein, D.F., Castellanos, F.X., Milham, M.P., 2010. Reliable intrinsic connectivity networks: test–retest evaluation using ICA and dual regression approach. Neuroimage 49 (3), 2163–2177.