NeuroImage 62 (2012) 891–901

Contents lists available at SciVerse ScienceDirect

NeuroImage

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

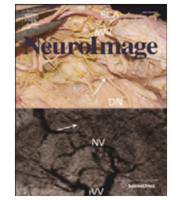

Review

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 . . . 891
(Spatial) ICA for FMRI . . . 892
Probabilistic ICA and MELODIC . . . 892
Loose methods ends . . . 895
Statistical inference . . . 895
Model order selection and the intrinsic dimensionality . . . 896
The 'killer applications' . . . 896
Artefacts and data denoising . . . 896
Data denoising . . . 897
Resting-state FMRI . . . 898
Future directions . . . 899
Variability and individual differences . . . 899
Multi-modal imaging . . . 899
Temporal dynamics and ICA . . . 899
Conclusion . . . 900
Acknowledgments . . . 900
References . . . 900

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
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.