Evaluation of ICA and Parallel ICA for Extracting Source Information from Simulated EEG-fMRI Signals

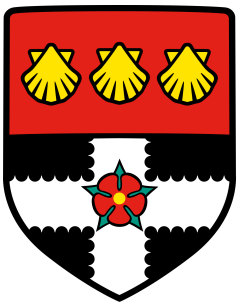

Asad Malik

School of Psychology and Clinical Language Sciences
University of Reading

This dissertation is submitted for the degree of
PhD in Cybernetics

August 2020

To Ammi and Abbu

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Asad Malik
August 2020

Acknowledgements

First of all, I would like to thank my primary supervisor, Etienne Roesch, who, mostly single-handedly, guided me through this PhD and has been very supportive throughout its course, not only in matters of science, but also in matters regarding my career development and personal life. I am also very grateful to my secondary supervisor, Kou Murayama, who joined this project in its final year and provided the feedback and assurances I needed to finalise this work. I am also thankful to the Engineering and Physical Sciences Research Council (EPSRC) for partially funding this work (EP/1549538 DTG 2014/2015).

My PhD experience would not have been the same without my lab members, especially Catriona Scrivener, Michael Lindner, and Ivano Ras, and I would like to thank them for the valuable discussions that contributed towards this work and for their friendship. I am also indebted to Slawomir Nasuto, who went out of his way to get me involved in EEG research and provided me the opportunity to acquire the experience that eventually landed me this PhD position.

I would also like to thank my colleagues at P1vital, especially Amy Bilderbeck, Andreea Raslescu, and Gerry Dawson, for keeping my workload balanced such that I always had sufficient time to work on my PhD, and for showing me how tasks can get done swiftly when part of an effective team.

I would not have started this journey without the financial and moral support of my parents, who encouraged me to take my time to explore what I wanted to do in life. These years will also be memorable as it is during them that I met my partner, Emma, who has been a constant source of inspiration and pragmatic wisdom.

Abstract

Parallel independent component analysis (ICA) is a framework for analysing concurrent electroencephalogram (EEG) and functional magnetic resonance imaging (fMRI) signals recorded from the brain that involves performing ICA in each modality and then matching the independent components (ICs) across modalities based on their statistical similarities. Together, the matched ICs are understood to provide information about the same neural sources (i.e. functional networks), with the EEG IC providing a high resolution temporal description and the fMRI IC providing a high resolution spatial description. In this thesis, EEG ICA, fMRI ICA, and parallel ICA are evaluated in terms of their accuracy at providing source information using synthetic data generated with The Virtual Brain (Sanz-Leon et al., 2013). Two novel extensions to parallel ICA, which are matching the ICs across modalities using spatial features and mutual information, are also proposed and evaluated.

The results of this work indicate that EEG ICA, fMRI ICA, and parallel ICA performances increase with the number of orthogonal sources in the absence of noise, and decrease with the level of noise dispersion when the number of sources are fixed. In the absence of noise, EEG and fMRI ICA performances do not vary largely with source network size (in regions), but in the presence of noise, they vary without clear trends. The incorporation of spatial features improves parallel ICA performance at matching the ICs across modalities, whereas the incorporation of mutual information, in comparison with correlation, deteriorates it. An important observation is that the single-modality and parallel ICAs do not always perform well in best-case conditions. That said, it must be acknowledged that this work is an initial investigation in this direction and further work with more diverse simulation parameters is needed to assess the generalisability of these results.

This thesis contributes to the existing body of ICA literature by performing the first evaluation of EEG ICA and fMRI ICA in terms of the number of neural sources and source network size. It is also the first evaluation of the parallel ICA approach that matches ICs across modalities using within-subject temporal dynamics, and the first application of parallel ICA that matches the ICs using spatial features and mutual information. This thesis is also the first demonstration of how The Virtual Brain can be used to evaluate unimodal and multimodal neuroimaging methods.

Table of contents

List of figures

List of tables

List of abbreviations

(p)IC	EEG IC, fMRI IC, and pIC
(p)ICA	EEG ICA, fMRI ICA, and pICA
BOLD	Blood Oxygenated Level Dependent
EEG	Electroencephalogram
EEG-fMRI	Concurrent EEG and fMRI
ERP	Event-related Potential
fMRI	Functional Magnetic Resonance Imaging
HRF	Haemodynamic Response Function
IC	Independent Component
ICA	Independent Component Analysis
IRF	Impulse Response Function
LFP	Local Field Potential
MI	Mutual Information
mm	Millimeter
ms	Millisecond
PCA	Principal Component Analysis
pIC	Parallel IC
pICA	Parallel ICA
PSP	Postsynaptic Potential
s	Second
sICA	Spatial ICA
SimTB	fMRI Simulation Toolbox
SNR	Signal-to-noise Ratio
tICA	Temporal ICA
TR	Repetition Time
TVB	The Virtual Brain

Chapter 1

Introduction

Since the first joint recording of electroencephalogram (EEG) and functional magnetic resonance imaging (fMRI) signals by Ives et al. (1993), several methods for their analysis have been developed. The motivation for concurrently recording signals from both modalities is to non-invasively measure human brain activity with the high temporal resolution of EEG and the high spatial resolution of fMRI. One framework for jointly analysing the two modalities is parallel independent component analysis (ICA), in which ICA is first applied to the signals of each modality separately to extract the underlying neural information they provide, and then combined across modalities so that, for some neural sources, there is an EEG description with high temporal resolution and an fMRI description with high spatial resolution. In order to better understand the type of neural information ICA and parallel ICA can provide and to characterise some of the conditions under which they can provide accurate information, in this thesis, both single-modality ICA and parallel ICA are evaluated using synthetic data. In addition to this, a basic, 'stripped-down' version of parallel ICA and two novel extensions to it are described and tested.

This chapter starts with a brief description of the physiological basis of EEG and fMRI signals in Section 1.1, followed by a statement on the motivation for combining the two modalities and some discussion on the extent to which they should be expected to provide information about the same neural sources in Section 1.2. A short overview of methods used to analyse concurrent EEG-fMRI signals is then presented in Section 1.3. Independent component analysis (ICA) is introduced in Section 1.4, followed by a description of parallel ICA and joint ICA in Section 1.5, which are two ICA-based methods for analysing EEG-fMRI signals. In Section 1.6 it is highlighted that, while ICA and ICA-based methods have been evaluated using synthetic data to some extent, the simulated source, EEG, and fMRI signals are typically not generated using a comprehensive pipeline of processes that model the propagation of biological activity at various scales. To this end, The Virtual Brain

(Sanz-Leon et al., 2013) is introduced as a multimodal simulator that produces synthetic data which is likely to be more biologically plausible than what has been previously used and, for this reason, is employed in this thesis. These sections provide the background to the aims of this project, which are described in Section 1.7 and, in Section 1.8, the structure of this thesis is described, along with a recommended path for its reading.

1.1 Physiological basis of EEG and fMRI signals

1.1.1 Electroencephalogram (EEG)

Brain activity involves the movements of ions which, at the cellular level, result in changes in the membrane potentials of neurons. This is the underlying basis of scalp potentials, which is what EEG electrodes measure. As explained in more detail below, these potentials are understood to be largely produced by the spatial and temporal summation of postsynaptic potentials (PSPs) of pyramidal cells in the cortex (Fisch and Spehlmann, 1999).

At rest, the difference between the intracellular and extracellular potential of a neuron is around 65 mV. Even at rest, there is a constant inflow and outflow of ions that occurs at a relatively slow rate. When the intracellular potential increases and reaches a certain threshold, gated ion channels open resulting in an influx of sodium ions followed by an efflux of potassium ions. This change in membrane potential is referred to as the *action potential*. This process is localised and usually starts near the axon hillock, where the threshold is the lowest (Fisch and Spehlmann, 1999), and then it propagates through the axon.

When the action potential reaches the axon terminal, it results in the release of a neurotransmitter at the synapse. Neurons that have receptors within the vicinity for that neurotransmitter temporarily open ion channels, resulting in a local change in their membrane potentials near the synapse (Fisch and Spehlmann, 1999). The *postsynaptic potential* (PSP) is the local potential of the membrane near a synapse in response to the activity of presynaptic neurons. PSPs can be excitatory or inhibitory. Excitatory PSPs increase the intracellular potential, making an action potential more likely, whereas inhibitory ones decrease it, making it more unlikely. EEG events can be due to both inhibitory and excitatory PSPs, and from an EEG signal we cannot tell which of the two resulted in an observed event (Fisch and Spehlmann, 1999).

EEG can only measure potential differences that result from currents large enough to be detected on the scalp surface. The current produced by an individual PSP therefore does not affect scalp potentials. To generate currents large enough to affect scalp potentials, populations of neurons need to have synchronised PSPs so that they produce large resultant

local field potentials (LFPs) that are the summations of the PSPs (Nunez and Silberstein, 2000). As pyramidal cells in the cortex are tightly arranged as columns and have fields with dipoles that have parallel orientation and similar timing, EEG is understood to largely measure summations of the LFPs they produce (Fisch and Spehlmann, 1999). Generally, only a small fraction of the currents in the cortex affect scalp potentials, as a significant portion of them are unable to penetrate between the various physical layers between the cortex and the scalp, such as cerebral spinal fluid and the skull (Fisch and Spehlmann, 1999).

EEG signals are typically measured as differences in the scalp potentials between pairs of electrodes. The electrodes most strongly detect electric fields that have dipoles that are perpendicular to the scalp surface and are generated in a large area of tissue located near them (Fisch and Spehlmann, 1999). As EEG is a direct measure of electrical activity, it has high temporal resolution and, typically, most EEG amplifiers can sample at 1000 Hz or more. In terms of localising brain activity, however, EEG has poor spatial resolution, as the signals are spatially mixed due to volume conduction and because only a small proportion of neural activity across the brain contributes towards the measurements.

1.1.2 Functional Magnetic Resonance Imaging (fMRI)

Unlike EEG, which is a direct measure of electrical activity in the brain, fMRI relies on changes in blood oxygen levels to make indirect inferences about underlying neural activations. The prevalent, somewhat simplified (Gjedde, 2001), view is that an increase in the neural activity in a brain area results in an increase in the cerebral blood flow to it to meet the energy demands of active neurons. This is the *haemodynamic response*; the fresh blood carries oxygen, in the form of oxyhaeomoglobin, along with other nutrients, such as glucose, to that part of brain to support the neural activity. This relationship between neural activity and blood flow is referred to as *neurovascular coupling*. Specifically, the generation and reception of neurotransmitters at synapses is thought to increase haemodynamic activity (Nunez and Silberstein, 2000).

In fMRI, the *blood oxygenated level dependent* (BOLD) signal is measured, which is considered to be an indicator of changes in deoxyhaemoglobin (Ogawa et al., 1990) and, therefore, of the haemodynamic response. It must be said, however, that it is still not well understood what it indicates in terms of neural activations and there are several factors, such as the strength of the magnetic field, acquisition parameters, and the thickness of blood vessels, that are known to affect it (Goense et al., 2010). Evidence that the BOLD signal does provide neural information primarily comes from studies that show that it is closely related to LFPs (Ashby, 2011; Logothetis, 2008).

Compared to EEG, fMRI has better spatial resolution, as it provides measurements at the voxel level, which typically have dimensions of 1 mm or more. The temporal resolution, however, of fMRI signals is poor. This is due to two reasons. The first is that fMRI signals have a slower sampling rate, with each volume typically being acquired in the order of seconds. The second is that information about underlying neural activations is temporally mixed as the haemodynamic response is a slow moving function that peaks in between 5 to 10 s after a neural area is activated and lasts around 20 s (Ashby, 2011; Goense et al., 2010). Therefore, by observing the haemodynamic response we cannot uniquely determine the time courses of the underlying neural events.

1.2 Motivation for combining EEG and fMRI signals

As EEG signals have high temporal resolution, but poor spatial resolution, and fMRI signals have high spatial resolution, but poor temporal resolution, the motivation for combining EEG and fMRI signals is to extract information about their shared neural generators with high spatiotemporal resolution, which cannot be obtained by either modality individually. The assumption behind this is that both signals contain some shared information about underlying neural sources. This is briefly discussed in this section.

If we consider a population of neurons within a voxel, then the electrical field of the voxel will only under certain conditions contribute towards scalp potential. Broadly speaking, as described in Section 1.1.1, some of these conditions are: the voxel should be close to the cortical surface, the dipole of the field should be perpendicular to the scalp, and several nearby dipoles should be parallel and synchronous. Therefore, it can be said that EEG and fMRI signals can contain shared neural information when the electrical fields of voxels meet these conditions and when the metabolic demands of the neural activity in the voxel are high enough to modulate the BOLD signal.

Nunez and Silberstein (2000) provide several examples of situations when EEG and fMRI signals would not contain shared neural information. For instance, they mention that the activities of stellate cells are not measured by EEG electrodes, as their electrical fields do not extend to the scalp, but they contribute significantly to BOLD, as they have high firing frequencies. In another example they describe that EEG alpha power is greatly reduced in the occipital areas when a participant opens their eyes, whereas the BOLD signal is not expected to attenuate as opening the eyes should not decrease visual processing. The reason they provide for this mismatch is that the EEG alpha oscillation observed in this case is the result of synchronisation between dipoles that is driven by a number of neurons that is too small to significantly modulate metabolic demand.

In their paper, Nunez and Silberstein (2000) describe forward models for both EEG and fMRI that map the local volume potentials of each voxel to both of the signals. The local volume potentials can comprise both sources and sinks, where the former are positive and the latter are negative. One evident difference between the two forward models is that, while both EEG and fMRI signals are functions of sources and sinks, EEG signals are always a function of the difference between sources and sinks within a local volume, whereas fMRI signals are not. This makes it explicitly clear that if there are similar numbers of positive and negative ions within local volumes, we will not see any corresponding effects in EEG signals, but we may see them in fMRI signals.

1.3 Methods for analysing EEG-fMRI signals

There are several methods for analysing concurrent EEG-fMRI signals. While the focus of this thesis is on ICA (see Section 1.4) and parallel ICA (see Section 1.5.1), a brief overview of these methods is provided here. This is not an exhaustive review, and a few of the most common approaches are only described. For more comprehensive reviews, see Abreu et al. (2018); Huster et al. (2012); Rosa et al. (2010).

Broadly speaking, EEG-fMRI analysis methods are classified as either being asymmetrical or symmetrical; in asymmetrical methods, information from one modality is used to 'constrain' the information acquired from the other modality and, in symmetrical methods, no explicit constraints that would 'prioritise' one modality over the other are placed. The two asymmetrical frameworks are (i) EEG-informed fMRI and (ii) fMRI-informed EEG, and they are described below. ICA-based methods fall into the symmetrical category, and they are described in Section 1.5.

EEG-informed fMRI In typical event-related fMRI analysis, a linear model is constructed for each voxel in which the dependent variable is the actual BOLD signal recorded at the voxel and the explanatory variables are artificial BOLD signals that simulate predicted responses to stimuli or events. These predicted BOLD signals are typically generated by convolving (see Section 1.5.1.2.1) stimulus time-courses, which are usually boxcar functions that are non-zero when the stimulus is present and zero otherwise, with a canonical haemodynamic response function (HRF). If the estimated coefficient (also referred to as the β-value) of an explanatory variable is statistically significantly not zero, then the activity at the voxel is understood to be modulated by the presence of the stimulus the variable corresponds to. In this manner, stimuli/events/conditions are correlated with voxels.

In EEG-informed fMRI, this linear model is typically extended to include explanatory variables that encode the predicted BOLD response to EEG features, such as event-related potential (ERP) amplitudes, ERP latencies, and band-power. Similar to standard fMRI analysis, the predicted BOLD signals corresponding to these features are generated by convolving their time-courses with a canonical HRF. Therefore, what is obtained from this analysis are voxels with activations that correlate with the encoded EEG features. For example, Novitskiy et al. (2011) used this method to identify brain regions that had activations that correlated with occipital P1 and N1 amplitudes, and Bénar et al. (2007) used it to identify regions with activations that correlated with P300 amplitude and latency. Similarly, Scheeringa et al. (2009) applied this method to EEG band-power time-courses to identify regions that had activity correlating with single-trial EEG alpha and theta power, and several studies have used it to investigate the neural correlates of resting-state alpha power (e.g. Gonçalves et al., 2006; Jann et al., 2009; Laufs et al., 2003).

fMRI-informed EEG As described in Section 1.1.1, the activities of populations of neurons in the brain produce electrical fields, which can be represented by dipoles and measured as local field potential (LFP) signals, and EEG signals are understood to be a linear mixtures of these LFP signals[1]. Therefore, one way to extract neural information from EEG signals is to attempt to identify the locations of the dipoles and their corresponding LFP signals. As biophysical forward models that map source (i.e. LFP) signals to scalp (i.e. EEG) signals have a many-to-one mapping, it is possible to use them to uniquely project source signals to scalp signals, but not possible to uniquely estimate source signals from scalp signals. Methods for EEG source localisation approach this *ill-posed inverse problem* by essentially estimating which configuration of sources is most likely to have generated the observed EEG under various constraints (Pascual-Marqui, 1999). While this problem cannot be solved using fMRI signals, it can be optimised by constraining the location of dipoles to voxels that are significantly modulated by experimental conditions. For example, Strobel et al. (2008) used this method to constrain estimated dipole locations of generators of the P3 in an auditory oddball task and Vanni et al. (2004) used it to localise visual ERPs.

Joint model inversion A symmetrical approach to EEG-fMRI analysis that has not been reviewed here but should be mentioned is the joint inversion of EEG and fMRI forward models. The approach is similar to EEG source localisation (see previous paragraph) with the addition that fMRI source localisation is simultaneously performed as well, i.e. sources are estimated by inverting the EEG and fMRI forward models such that they are most likely to

[1] Ignoring non-brain artefacts, such as power line noise.

have generated the observed EEG-fMRI signals. This approach has been formally described in Valdes-Sosa et al. (2009) and has been applied in various ways, e.g. in Wei et al. (2020) the source estimation is done within the framework of dynamic causal models (Friston et al., 2003).

1.4 Independent Component Analysis (ICA)

Independent component analysis (ICA) is a method for decomposing a matrix, Y, into the product of a mixing matrix, M, and a source matrix, S, i.e.:

$$Y = MS. \tag{1.1}$$

ICA estimates M and S such that each row of S, or independent component (IC), is maximally statistically independent[2] from the others. The basic idea is that if we can assume that Y is a linear mixture of statistically independent, latent, sources, then ICA provides a means for obtaining source information (i.e. S) from the mixture (i.e. Y). M is referred to as the mixing matrix as it contains the weights with which the ICs are linearly mixed in Y.

When applied on EEG or fMRI signals, the ICA can either be temporal or spatial based on how Y is structured. For temporal ICA, the columns of Y are different time points and the rows are different spatial points (see Figure 1.1), whereas for spatial ICA, the rows are different time points and the columns are different spatial points (see Figure 1.2). Correspondingly, as ICA optimises for independence between the rows of S, in temporal ICA, S is estimated such that ICs are maximally independent from each other in the temporal domain, whereas in spatial ICA they are estimated to be maximally independent in the spatial domain. Also, based on the algebraic formulation, in temporal ICA, the ICs are time-courses and their corresponding columns in the mixing matrix are their spatial maps, whereas in spatial ICA, the ICs are spatial maps, and their corresponding columns in the mixing matrix are their time-courses. Typically, when the number of time points is less than the number of spatial points, spatial ICA is chosen, and when the number of spatial points is less than the number of time points, temporal ICA is chosen (Calhoun et al., 2001c). Therefore, typically (and in this thesis), temporal ICA is performed on EEG signals and spatial ICA on fMRI signals.

One other form of ICA that is mentioned in this thesis (see Section 1.5.1.2.4), but not applied, is spatiospectral ICA. Spatiospectral ICA is a variant of spatial ICA that has been

[2]Two variables are statistically independent when the product of their marginal probability density functions is equal to their joint probability density function. Intuitively, what this means is that when two variables are independent, information of one variable does not tell you anything about the other.

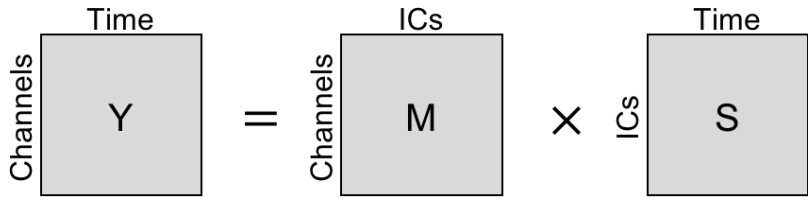

Figure 1.1 EEG temporal ICA. The data (Y) is structured such that the rows are channels and the columns are time-points. Correspondingly, the ICs, i.e. rows of the source matrix (S), are time-courses, and their corresponding columns in the mixing matrix (M) are spatial (scalp) maps.

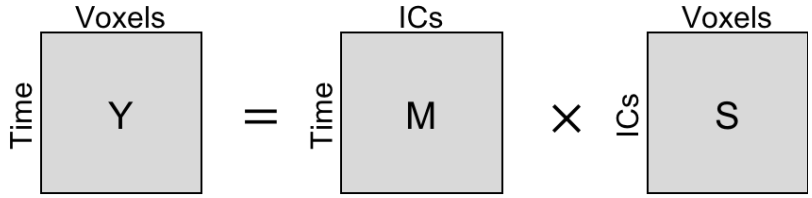

Figure 1.2 fMRI spatial ICA. The data (Y) is structured such that the rows are time-points and the columns are voxels. Correspondingly, the ICs, i.e. rows of the source matrix (S), are spatial maps, and their corresponding columns in the mixing matrix (M) are time-courses.

applied on EEG signals (Bridwell et al., 2016, 2013; Wu et al., 2010). The difference between spatiospectral ICA and spatial ICA is that ICA is not performed on the raw signals but rather on their spectrograms. That is, the activity recorded at each channel is first decomposed into different frequency bands of interest, and then Y is constructed such that there is a column for each channel and frequency band combination that contains the mean spectral power within time windows (or epochs), which are indexed by the rows of Y. Consequently, each IC (row of S) comprises a spatial map for each power band, and the corresponding column of M is its time-course.

The estimation of M and S such that the rows of S are statistically independent from each other is a nontrivial mathematical problem for two main reasons. The first is that for a Y there may not exist an S with statistically independent rows. The second is that estimating S requires calculating the exact marginal and joint probability densities of its rows, which is computationally expensive. Therefore, ICA algorithms estimate the rows of S such that they are *maximally*, not exactly, statistically independent from each other, using a variety of heuristics. Throughout this thesis, we have used FastICA (Hyvärinen and Oja, 1997) with the tanh nonlinearity (see Section 2.2.3.1 for further description).

ICA is often termed as being a 'model-free' approach towards the analysis of neural signals as it does not use any biological information on how the signals were generated. There

are a few assumptions, however, that it makes. One of these is that the signals are assumed to be linear mixtures of the sources, as ICA decomposes the signals into linear combinations of ICs. Another is that ICA assumes that sources are statistically independent of each other, again as it decomposes the signals into ICs that are maximally statistically independent. The third assumption is that at most only one source is Gaussian, as ICA can only separate ICs that are non-Gaussian (Hyvärinen and Oja, 2000). In fact, most ICA algorithms, e.g. FastICA (Hyvärinen and Oja, 1997), estimate ICs by optimising for non-Gaussianity (see Section 2.2.3.1).

ICA can be extended to work with the data of multiple subjects using techniques such as concatenation (Calhoun et al., 2001b) and tensor algebra (Beckmann and Smith, 2005). In this thesis, group ICA has not been used. However, the interested reader may look at Appendix D in which the temporal concatenation approach (Calhoun et al., 2001b) that is used in the methods described in Section 1.5 is explained.

1.5 ICA-based methods for analysing EEG-fMRI signals

As ICA is relatively model-free and data-driven, it can be used for the symmetric analysis of joint EEG-fMRI signals without having to make the biological assumptions that are needed for forward models in joint model inversion methods (see Section 1.3). The two most commonly used methods for EEG-fMRI analysis using ICA are parallel ICA and joint ICA, and both are described in this section. In this thesis, the focus is on the evaluation and development of parallel ICA, as it provides a flexible framework for integrating the two modalities that, unlike joint ICA, allows some margin for the sources to be mixed differently across modalities (explained in Section 1.5.2.1) and generally enables the manipulation of several parameters that could be experimented with to improve performance (more details in Section 1.5.1.2.5). There are other ICA-based methods that could be used for EEG-fMRI, such as Linked ICA (Groves et al., 2011), that have not been described in this section for the sake of brevity and because they have not significantly informed the work presented in this thesis.

1.5.1 Parallel ICA

As mentioned in the introduction of this chapter, parallel ICA for EEG-fMRI involves performing ICA on each modality and then matching the ICs across modalities based on their statistical similarities. Together, for a source, the matched ICs provide multimodal information, with the EEG IC providing a high resolution temporal description and the fMRI

IC providing a high resolution spatial description. Parallel ICA algorithms can be divided into two groups: those that perform 'matching across subjects' and those than performing matching on 'trial-by-trial' data. These are described in the following subsections.

1.5.1.1 Matching across subjects

Parallel ICA was first proposed and performed by Liu and Calhoun (2007) using the 'matching across subjects' approach. In this method, for each modality, single-subject analysis is first performed and, for each subject, a summary endpoint is obtained, e.g. an EEG ERP time-course or an fMRI statistical parametric map. For each modality, each subject's endpoint is arranged as a row vector and the 'group' data matrix, Y (see Section 1.4), is constructed by vertically concatenating the endpoints across subjects. ICA is then performed on Y for each modality and the ICs are matched across modalities if their corresponding columns in the mixing matrices correlate. If they do correlate, this means that the ICs are mixed similarly across subjects, suggesting that they correspond to the same source. One additional feature of this method is that the ICA decompositions across modalities can interact. For this, a modified Infomax (Bell and Sejnowski, 1995) algorithm is used in which the optimisation is not only based on finding independent components, but also on finding a pair of components across modalities that have high correlation[3].

In Liu and Calhoun (2007), this method was applied to EEG-fMRI data collected from an auditory odd-ball task. Each row of Y_E (the EEG data matrix) was the averaged ERP for a target tone condition recorded from Cz for each subject, and each row of Y_F (the fMRI data matrix) was the statistical parametric map of each subject for the same condition. One pair of highly correlated components ($r = 0.74$) was reported (see Figure 1.3) in which the fMRI IC spatial map had high activations in the superior and middle temporal gyri and the EEG IC time-course had the N100 peak. This method can be extended to more than two modalities, as was shown in Vergara et al. (2014) where it was applied on structural MRI, fMRI, and genetic data. While this three-way approach has not as yet been applied on EEG signals, there do not appear to be any clear reasons why that could not be done.

The 'matching across subjects' parallel ICA approach has been evaluated using simulated signals. In Liu and Calhoun (2007), eight sources were constructed and mixed across twenty-five subjects in an fMRI matrix Y_F (25 x 39800) and an EEG matrix Y_E (25 x 4510). For each modality, a mixing matrix was pseudo-randomly generated, such that a pair of columns (one from each modality) was correlated with each other. The strength of the correlation

[3] An implementation of this algorithm is provided in the Fusion Toolbox (Rachakonda et al., 2008) and it provides users with the option to either maximise the correlation between a column of the mixing matrix of one modality with a column in the other, or between a row of the source matrix of one modality with a column of the mixing matrix of the other.

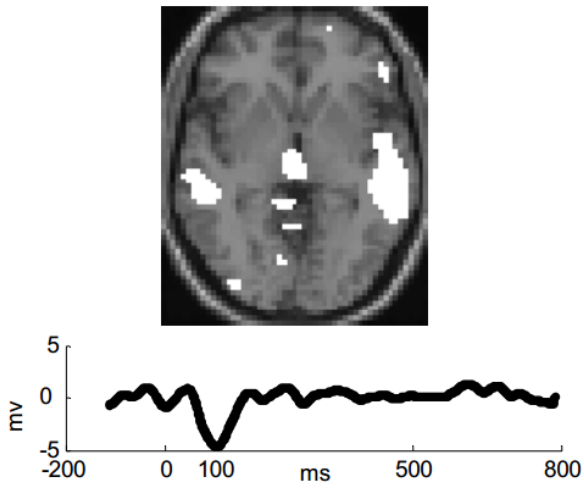

Figure 1.3 A parallel IC (fMRI IC spatial map and EEG IC time-course) using the matching across subjects approach. The EEG IC and fMRI IC are interpreted as corresponding to the same source. Reprinted from Liu and Calhoun (2007) with permission. Copyright © 2007, IEEE.

was used as the ground-truth connection strength and Gaussian noise was also added to the mixed signals. ICA performance was evaluated in terms of component accuracy (accuracy of the components corresponding to sources within each modality) and connection accuracy (accuracy of the connection strength for the pair of components corresponding to the same source). The parameters that were varied were: the connection strength, the signal-to-noise ratio (SNR), and the number of ICs. A comparison was also performed between 'interactive' ICA (which had a bias towards increasing the correlation between a pair of ICs across modalities, see above) and standard ICA (each modality's ICA was performed independently of the other). Overall, both methods had high component accuracy (> 0.9), when the number of ICs was eight or more, and the 'interactive' method had slightly better connection accuracy, except when the ground-truth connection strength was less than 0.5, in which case it was inflated. Evaluation using simulated signals was also performed for the three-way method in Vergara et al. (2014). The authors observed that three-way parallel ICA performed more accurately than parallel ICA between two modalities and ICA within single modalities, and that its performance improved the stronger the connection strength between the modalities was. Three-way parallel ICA was more accurate than the other two methods at estimating the connection strength as well.

1.5.1.2 Trial-by-trial integration

There are two main differences between the 'trial-by-trial' parallel ICA methods and the 'matching across subjects' parallel ICA methods. The first one is that the latter match ICs across modalities based on how similarly they are mixed across subjects whereas the former match the ICs based on the similarity of their trial-by-trial variation. The second one is that in the 'trial-by-trial' methods the ICA decomposition in each modality is independent of the decomposition in the other modalities, whereas in the 'matching across subjects' approach, they can interact, as described in the previous section.

Two methods for 'trial-by-trial' parallel ICA are described in some detail in the following subsections as they provide the main background for the work in Chapter 4. This is followed by a short description of two methods that use spatiospectral ICA. A basic, 'stripped-down' version of parallel ICA and some possible extensions to it are then described. Convolution and deconvolution are briefly explained first as they are used by the methods.

1.5.1.2.1 Convolution and deconvolution Convolution is a mathematical operation using which the output of a linear, time-invariant system for a given input signal can be predicted if its *impulse response function* (IRF) is known. Functional MRI signals are usually assumed to be outputs of such a system, e.g. in classical fMRI analysis, statistical parametric maps are obtained by constructing a predicted fMRI signal by convolving the stimulus time-course (the input signal) with a haemodynamic response function (the IRF) and identifying voxels that significantly correlate with it. Deconvolution refers to the inverse of this process, i.e. instead of constructing the predicted output signal, a signal is decomposed into either the estimated input signal it is in response to (if the IRF is known) or its IRF (if the input signal is known).

Convolution can be used to match EEG and fMRI signals in the temporal domain. This can be done by first convolving the EEG signal with a canonical haemodynamic response function (HRF) to construct a predicted fMRI signal and then identifying fMRI voxels that correlate with it. This approach assumes a shape of the HRF and that it is invariant across different spatial locations in the brain and subjects. The same process can be applied to other time-courses extracted from EEG and fMRI signals, such as their IC time-courses (e.g. Eichele et al., 2008, more details below). Similarly, deconvolution can be used to match EEG and fMRI signals in the temporal domain. This can be done by first estimating the IRF by deconvolving the fMRI signal using the EEG signal as the input signal and then checking if the IRF has properties of an HRF. This process can be applied to IC time-courses as well (e.g. see Figure 1.4). In the sections that follow, four different 'trial-by-trial' parallel ICA algorithms are described.

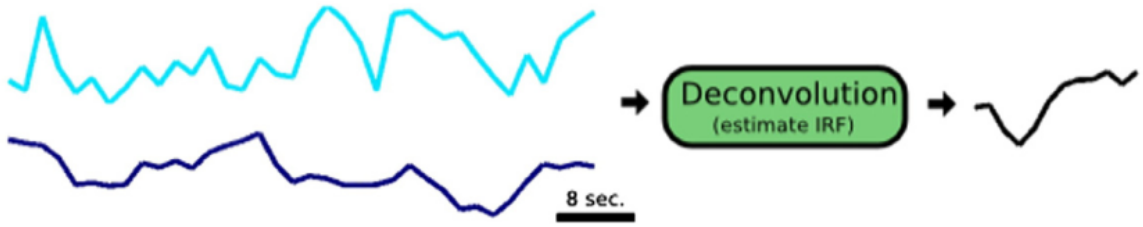

Figure 1.4 Estimating the impulse response function (IRF) by deconvolving an fMRI IC time-course (dark blue) with an EEG IC time-course (light blue) (see Section 1.5.1.2.1). The estimated IRF (black) is shown on the right. Adapted from Bridwell et al. (2013) with permission. Copyright © 2012, Elsevier.

1.5.1.2.2 Method 1: Eichele et al. (2008) The basic idea in the method proposed by Eichele et al. (2008) is to match the EEG ICs and fMRI ICs with each other using linear regression between their single-trial weights (defined below). The first step in this method is that group ICA (see Appendix D) is performed in each modality separately and then single-subject ICs and mixing matrices are obtained using back-reconstruction (see Section D.5). The ICs are then evaluated against various statistical criteria[4] and those that do not meet the criteria are excluded from further analysis. To match the EEG and fMRI ICs in the temporal domain, they are co-registered at the trial-level using single-trial weights. For the fMRI ICs, this process is trivial, as their single-trial weights are simply their values at each time point as, in the paradigm described in Eichele et al. (2008), a single volume is acquired per trial. For the EEG ICs, however, the process is not as straightforward as the EEG signals are acquired at a much higher sampling rate. In this process, for each subject, for each trial, the EEG signals across channels are concatenated into a single vector and, similarly, the back-projected time-courses of each IC (see Section D.6) are also concatenated across channels into their individual vectors. Then, a linear model is fit with the EEG vector as the dependent variable and the IC vectors at the explanatory variables. For each IC, the single-trial weight is its estimated β-value (coefficient).

For each subject, once the EEG and fMRI IC single-trial weights have been obtained, the EEG and fMRI ICs are matched by fitting a linear model between them across trials. For

[4]The details of these criteria have not been described here as, to a degree, they are task-specific. See Eichele et al. (2008) for further details.

each fMRI IC, the following model[5,6] is estimated:

$$f_i(t) = \beta_s s(t) + \beta_1 e_1(t) + \cdots + \beta_n e_n(t) + \beta_0, \tag{1.2}$$

where t indexes trials, $f_i(t)$ are the single-trial weights (i.e. timecourse) of the i-th fMRI IC, $e_k(t)$ are the single-trial weights of the k-th EEG IC convolved with the HRF, $s(t)$ is an on-off function encoding the presentation of the stimulus convolved with the HRF, and n is the number of EEG ICs. Once the β-values for each EEG-fMRI IC pair have been obtained for each subject, they are tested for statistical significance across subjects. EEG ICs and fMRI ICs that have β-values that are statistically significantly non-zero are matched.

This method was used on data collected from an auditory oddball paradigm which had standard (frequent) tones and target (infrequent) tones and the results were reported for single matching EEG-fMRI IC pair (see Figure 1.5). The EEG IC time-course was significantly different for the standard and target conditions at multiple time points between 100 and 350 ms after stimulus onset. The fMRI IC spatial map had significant clusters of activation in the superior temporal gyri, temporal poles and the anterior cingulate gyrus. These results were considered plausible, as they included areas known to be involved in auditory processing, sensory discrimination, and novelty-related functions.

1.5.1.2.3 Method 2: Eichele and Calhoun (2010) The basic idea of the method proposed by Eichele and Calhoun (2010) is the same as that of Eichele et al. (2008) in that the EEG and fMRI ICs are matched using linear regression between their single-trial weights. The main difference is in how the fMRI IC single-trial weights are calculated[7]; in Eichele et al. (2008), the fMRI IC time-courses are used as their single-trial weights, as a single fMRI volume is acquired per trial, whereas here they are extracted using deconvolution and regression. The first step for this is that, for each fMRI IC, the IC time-course is deconvolved with respect to the stimulus time-course to estimate the IRF contained in the IC in response to the stimulus (see Section 1.5.1.2.1). Then, for each selected[8] fMRI IC, for each trial, a predicted BOLD signal is constructed by convolving the onset of the trial with the IRF of the IC. A linear model is then fit with the dependent variable being the fMRI IC time-course and

[5] In Eichele et al. (2008), the EEG single-trial weights were detrended and the $e_k(t)$ were orthogonalised with respect to $s(t)$, for each k, before model fitting. These steps have not been described in detail as they are task-specific. See Eichele et al. (2008) for further details.

[6] It is assumed that the intercept, β_0, was added. This has not been mentioned in Eichele et al. (2008).

[7] The EEG IC single-trial weights in Eichele and Calhoun (2010) were also not calculated the same way as they were in Eichele et al. (2008). However, for these, the method was much more task-specific, and, therefore, has not been described here.

[8] In Eichele and Calhoun (2010), fMRI ICs that had IRFs that resembled HRFs were selected (see Section 1.5.1.2.1).

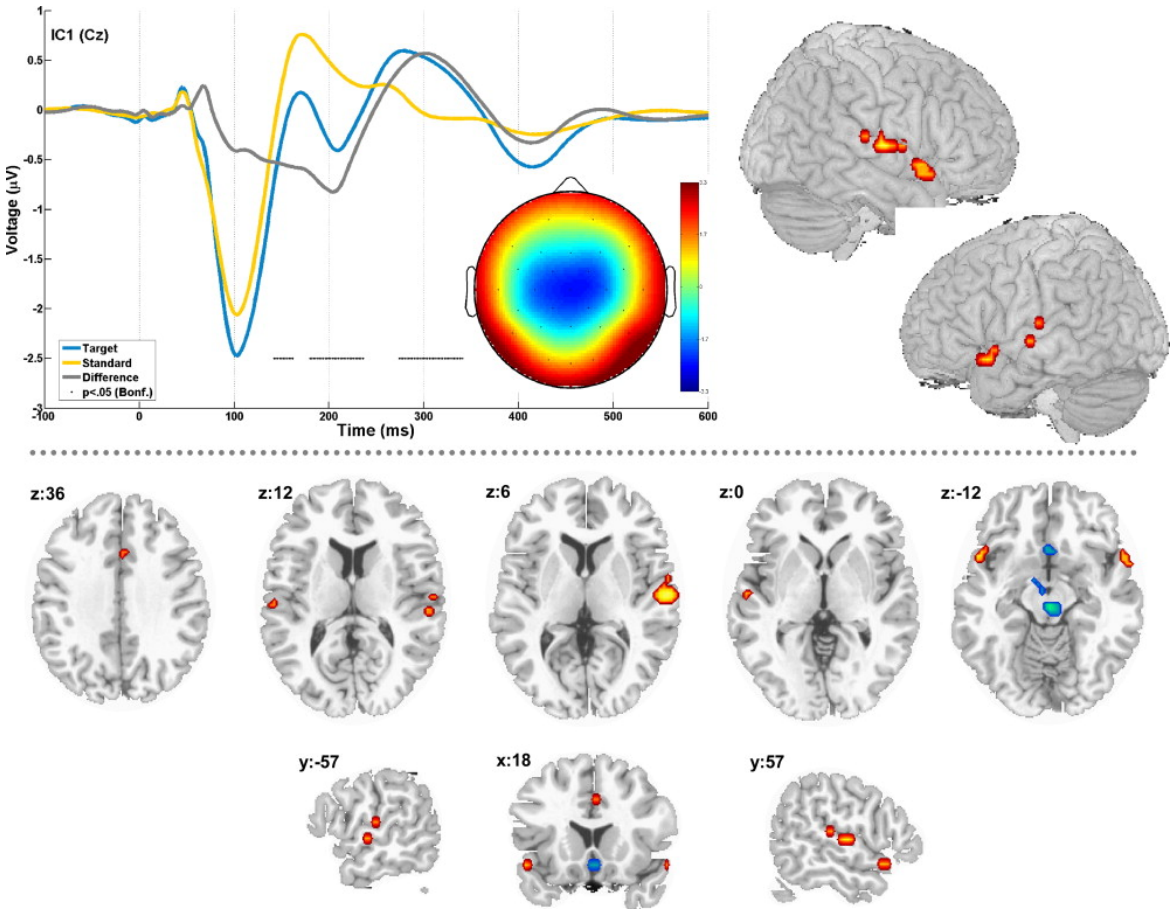

Figure 1.5 A parallel IC using the trial-by-trial approach described in Section 1.5.1.2.2. Top left: The average back-projected EEG IC time-course on channel Cz is shown for the standard (yellow) and target (blue) conditions. The difference between the two time-courses is shown in grey, and the dots indicate where it is statistically significant. The EEG IC scalp map also shown. Top right and bottom: The thresholded fMRI IC spatial map. Reprinted from Eichele et al. (2008) with permission. Copyright © 2008, Elsevier.

the explanatory variables being the predicted BOLD signals of the trials. For each trial, the estimated β-value corresponding to it is used as its single-trial weight for the fMRI IC.

Once the EEG and fMRI IC single-trial weights have been obtained[9], they are matched using the following pairwise linear models[10,11]:

$$e_i(t) = \beta_1 m(t) + \beta_2 f_j(t) + \beta_0, \tag{1.3}$$

where t indexes trials, $e_i(t)$ is single-trial weight of the i-th EEG IC, $f_j(t)$ is the single-trial weight of the j-th fMRI IC, and $m(t)$ is the mean of the single-trial weights across all the fMRI ICs. Similar to Eichele et al. (2008), EEG and fMRI ICs that have β-values that are statistically significantly non-zero across subjects are matched.

1.5.1.2.4 Spatiospectral ICA methods Two methods proposed by Bridwell et al. (2013) and by Wu et al. (2010) use spatiospectral ICA (see Section 1.4) on the EEG signals, instead of temporal ICA, and are applied on resting state data[12]. As the evaluation and development of parallel ICA in this thesis (Chapter 4) is not based on these methods, they are not described in detail here.

The basic idea of the method in Bridwell et al. (2013) is that EEG and fMRI ICs are matched based on the inspection of the IRFs obtained through deconvolution (see Section 1.5.1.2.1). Spatiospectral ICA is performed on the spectrograms of the EEG signals and, in order to have the EEG IC time-courses and fMRI IC time-courses co-registered in time, the windows of the spectrograms (or epochs) are set to have the same as length as the acquisition time of fMRI volumes (i.e. TR). For each pair of selected EEG and fMRI ICs, the fMRI IC time-course is deconvolved with respect to the EEG IC time-course to obtain an estimated IRF (see Figure 1.4). The ICs are then matched across modalities if their IRFs looked similar to biologically plausible HRFs by testing if they have non-zero values at 4, 6, 8, or 10 s. In terms of interpretation, for a source, fMRI ICs provide spatial maps and the matching EEG ICs provide spectral time-courses.

The method used in Wu et al. (2010) is qualitatively different from the previous ones in that no process is employed to decide which EEG and fMRI ICs match each another. Rather,

[9] Similar to Eichele et al. (2008), some EEG and fMRI ICs were excluded from the matching process based on various statistical criteria. See Eichele and Calhoun (2010) for further details.

[10] In Eichele and Calhoun (2010), $f_j(t)$ was orthogonalised with respect to $m(t)$ before model fitting, and the same analysis was performed with the differentials of the EEG and fMRI IC single-trial weights.

[11] It is assumed that an intercept, β_0, was added. This has not been mentioned in Eichele and Calhoun (2010).

[12] As these methods are applied on resting state data, strictly speaking, they are not trial-by-trial, as there are no trials. For the purposes of this thesis, however, they have been classified as 'trial-by-trial' as within-subject temporal variation is used to match the ICs across modalities.

a single EEG IC and a set of fMRI ICs are selected based on properties that suggest their relevance to the research question[13], and then all IRFs extracted from the fMRI ICs using deconvolution with respect to the EEG IC are considered to provide relevant information on how the activations in various parts of the brain are related to the temporal dynamics of the EEG IC.

1.5.1.2.5 'Stripped-down' method and extensions The 'trial-by-trial' parallel ICA methods for task data described above involve group analysis and temporal reduction using single-trial weights. The performance of the simplest, 'stripped-down' version of parallel ICA, which does not involve either of these steps, has never been demonstrated, even though, arguably, this is the cornerstone of the approach and should be used as the benchmark against which any extensions, such as using single-trial weights rather than time-courses, should be compared. Group ICA, for example, adds layers of complexity through concatenation, dimensionality reduction, and back-reconstruction (see Appendix D), all of which will, to a degree, reduce the likelihood of identifying individual-specific effects. In our opinion, the process of method development should involve first clearly demonstrating and reporting how the method performs at the single-subject level, even if it is evident that the effects are too small to be detected, before group data is used.

The same argument applies to reducing the temporal information in the data by using single-trial weights rather than signal or IC time-courses. Again, in our opinion, in the first instance of investigating parallel ICA, the ICA in each modality should be performed on the signals of that modality that have not been temporally reduced, even if the reduction may improve the matching of ICs across modalities for particular datasets.

Apart from the fact that 'trial-by-trial' parallel ICA has not been demonstrated in its 'stripped-down' form, there are a few other limitations to how it has been used. One of these is that the matching of ICs across modalities is only performed in the temporal domain by either convolving the EEG IC time-courses (or single-trial weights) with a canonical HRF or by deconvolving the fMRI IC time-courses to extract an HRF. Potentially, the matching could also be performed in the spatial domain, e.g. by projecting fMRI IC spatial maps to EEG IC scalp maps. Therefore, one possible extension to parallel ICA is matching the ICs across modalities using their spatial features.

Another limitation of the existing methods is that the ICs are matched across modalities based on their linear association, e.g. in both Eichele et al. (2008) and Eichele and Calhoun (2010), linear regression is used to measure the similarity between the EEG and fMRI ICs.

[13] In Wu et al. (2010), the authors were interested in the brain activations underlying EEG alpha oscillations and therefore selected an EEG IC with high alpha power and fMRI ICs that had activations reported to be linked with alpha power in previous literature.

Therefore, other possible extensions to parallel ICA are matching the ICs using nonlinear measures, such as mutual information and transfer entropy (see Appendix E).

As is described in more detail in Section 1.7, in this project, the 'stripped-down' parallel ICA method, which uses the data of a single-subject and does not perform any temporal reduction on the data pre- or post-ICA in either modality, is demonstrated and evaluated, and it is extended to use spatial features and mutual information for matching the ICs.

1.5.2 Joint ICA

Joint ICA for EEG and fMRI fusion was first proposed by Calhoun et al. (2006b). The main difference between joint ICA and parallel ICA is that in parallel ICA, ICA is performed in each modality separately and then the ICs are matched across modalities, whereas in joint ICA, data from both modalities is concatenated into a single matrix and then ICA is performed on that matrix. Like parallel ICA, joint ICA methods can also be grouped in 'matching across subjects' methods and 'trial-by-trial' methods.

1.5.2.1 Matching across subjects

Similar to 'matching across subjects' parallel ICA (see Section 1.5.1.1), this version of joint ICA also uses single-subject summary endpoints. For example, in EEG-fMRI joint ICA, for each subject, a row vector is constructed by concatenating an fMRI spatial map and an ERP time-course that correspond to the same event, and then the row vectors of all subjects are vertically concatenated to form the data matrix (Y) on which ICA is performed (see Section 1.4). The ICs obtained, therefore, also contain an fMRI spatial map and an ERP time-course, which, by virtue of being in the same IC, are interpreted as matching the same source. Joint ICA, therefore, provides a means to split these EEG and fMRI features into smaller components based on how they co-vary across subjects. For example, different ICs may split an ERP into its peaks, such as the C1, N1, and P1. Each of these peaks would then, as a result of the ICA, have an associated fMRI spatial map. As, in joint ICA, a single IC comprises both the EEG and fMRI information, joint ICA can only extract information on sources that are mixed exactly the same way in both modalities across subjects. This is in contrast to 'matching across subjects' parallel ICA, which allows some margin for differences between how the sources are mixed across subjects in the two modalities as matching ICs are not required to be perfectly correlated (Calhoun et al., 2009).

In Edwards et al. (2012), joint ICA was used to combine ERPs and fMRI spatial maps obtained in a go/no-go task and in Calhoun et al. (2006b) it was used to combine them in an auditory oddball paradigm. Joint ICA has also been used to combine other modalities. For

example, in Calhoun et al. (2006a), it was used to combine structural MRI with fMRI. One of the benefits of joint ICA over single-modality ICA is that it makes it more likely to get information about sources that have contributions to the signals in both modalities that are too small within each modality to have been identified on their own. For example, in Calhoun et al. (2006b), the authors identify a peak in the P3 that may have some correspondence with brainstem activity. An analysis on the EEG signals alone would have been unlikely to uncover this, given the small contributions the brain stem is expected to make to EEG signals because of its distance from the scalp (see Section 1.1.1).

1.5.2.2 Trial-by-trial integration

A method for applying joint ICA on within-subject trial-by-trial data is presented in Moosmann et al. (2008). In this method, for each subject, for each trial, a column vector is created by concatenating the fMRI spatial map (each trial is a single fMRI volume) with the EEG signals of all the channels recorded during the acquisition of the volume. To bring them to the same temporal scale, either the EEG signals are convolved with a canonical HRF, or the fMRI signals are deconvolved using a canonical HRF as an IRF to extract the underlying input signal (see Section 1.5.1.2.1). Group ICA (see Appendix D) is then performed on a matrix in which the data for multiple subjects is horizontally concatenated and then subject-specific joint ICs are obtained using back-reconstruction (see Section D.5).

Unlike the 'trial-by-trial' parallel ICA methods, this method has been evaluated using synthetic data in Moosmann et al. (2008). For the evaluation, signals of 24 subjects were simulated, with six sources that had hypergaussian temporal distributions. Each source had an fMRI time-course (across trials), an fMRI spatial map, an ERP time-course (within a trial) and an EEG scalp map. The ERP time-courses were constructed using actual EEG signals. Only four of the six sources were selected for each subject randomly, and noise was added to increase the between-subject variance. Twenty ICs were estimated and back-reconstructed (see Section D.5) to single-subject space. The similarity of each IC with the sources was measured using regression: for each source feature (fMRI time-course, fMRI spatial map, ERP time-course, and EEG scalp map) and IC, regression was performed with the source feature as the dependent variable and the corresponding IC feature, averaged across subjects, as the explanatory variable and, for each source, the sum of the fit (R^2) of the four models was used as the similarity measure. For the six sources, the sum of the fits of the ICs with highest similarity were: 378%, 310%, 290%, 366%, 349%, and 306%, and these were considered to be good results.

1.6 Evaluation of ICA and ICA-based EEG-fMRI analysis methods

Generally, ICA provides a simple and elegant method for separating sources that are linearly mixed based on their statistical independence. In the cases of both EEG and fMRI signals, using ICA to uncover information on latent neural activity appears to be plausible as networks performing different functions may be expected to behave independently, have non-Gaussian distributions, and be linearly mixed[14] (Makeig et al., 1996; Vigario et al., 2000). However, as is the case with several methods used in human neuroscience, rigourous testing has not been performed to evaluate the extent to which ICA and ICA-based EEG-fMRI methods actually provide accurate biological information; the testing that has been done is typically limited to checking the proof-of-concept of novel methods (see below). The primary constraint that makes the evaluation of these methods difficult is that real, biological source signals (e.g. LFPs) are impossible to non-invasively obtain and, therefore, are rarely readily available for such investigations. The compromise that has to be made, therefore, is that synthetic data, constructed from mathematical models depicting our somewhat 'simplified' understanding of biological processes, has to be used.

The evaluation of ICA and ICA-based methods for combining EEG-fMRI signals using simulated signals is not a new concept and has been previously performed in several cases (e.g. Bridwell et al., 2016; Ghahremani et al., 1996; Liu and Calhoun, 2007; Moosmann et al., 2008). However, in most of these cases, the EEG and fMRI signals have been generated using 'in-house' methods that do not model biological processes comprehensively. In this thesis, source (LFP), EEG, and fMRI signals are simulated using The Virtual Brain (TVB) (Sanz-Leon et al., 2013). The primary reason for using this simulator is that, to our knowledge, this is the only simulator that generates concurrent EEG and fMRI signals. The secondary reason is that TVB, compared to other simulators, in our opinion, is more biologically plausible as it implements a pipeline of processes that model the propagation of biological signals at multiple temporal and spatial scales while accounting for local and long-range connectivity (see Section 2.2.1 for more details).

Various methods have been used for the construction of source time-courses for the simulation of EEG signals in existing literature. For example, in Moosmann et al. (2008) and Lei et al. (2010), the source time-courses were constructed by combining ERPs, extracted from real EEG signals, with stimuli time-courses, which varied in their spectral properties across sources; in Bridwell et al. (2016), source time-courses of resting state EEG were constructed using wavelet decompositions of real resting state EEG signals; and in Ghahremani et al.

[14] These are the assumptions ICA makes, see Section 1.4.

(1996), the source time-courses were audio signals. Various methods have also been used in projecting these source time-courses to the scalp to get EEG signals. In some of these methods, the projection is performed by placing the sources at fixed anatomical locations and projecting their signals to the scalp using head models (e.g. Ghahremani et al., 1996; Lei et al., 2010). In other cases, the source scalp maps are constructed directly without any forward modelling (e.g. Bridwell et al., 2016; Moosmann et al., 2008). In these methods, typically the noise is drawn from Gaussian distributions (e.g. Bridwell et al., 2016; Lei et al., 2010), but uniform distributions have also been used (e.g. Ghahremani et al., 1996).

The most common method for simulating an fMRI signal is to convolve the stimulus time-course with an HRF (e.g. Lei et al., 2010), but other methods, such as using time-courses with various spectral properties (Moosmann et al., 2008) or just using on-off signals (Calhoun et al., 2001c), have also been used. One fMRI simulator that should be mentioned is SimTB (Erhardt et al., 2012), which was used in the evaluation of 'three-way' parallel ICA (Vergara et al., 2014, see Section 1.5.1.1). Compared to TVB, this simulator has some additional features: it provides source templates with spatial maps that are similar to functional networks observed in real fMRI, making it easier to construct sources with biologically plausible anatomy, and it enables the addition of typical fMRI artefacts, such as head motion, and the addition of sensor noise, making the properties of the fMRI signals more realistic. However, SimTB only models brain activity at the temporal scale of fMRI signals, whereas TVB operates at the LFP-level of cortical vertices and, therefore, can model the dynamical properties of neural signals with much higher temporal resolution and, consequently, can provide the user with a selection of neural mass models which determine the shape of the evoked LFP responses. In addition to this, TVB implements local connectivity, which is instantaneous and within the neighbourhood of a cortical vertex (see Section 2.2.1), and long-range connectivity, which is between regions and time-delayed, and provides various parameters for simulating source noise, which are all features that SimTB does not provide. Most importantly, TVB provides multiple monitors for measuring brain activity for different modalities concurrently, such as EEG, fMRI, and MEG, whereas SimTB only simulates fMRI signals.

1.7 Aims

The aims of this project are listed below, and their context is summarised in the text that follows.

1. To evaluate EEG temporal ICA and fMRI spatial ICA in terms of their accuracy at providing information about functional networks in the brain using signals simulated with TVB.

2. To evaluate a 'stripped-down' version of 'trial-by-trial' parallel ICA that works on the data of a single-subject and does not reduce the EEG or fMRI time-courses to single-trial weights using signals simulated with TVB.

3. To extend parallel ICA by matching the ICs across modalities using spatial features.

4. To extend parallel ICA by matching the ICs across modalities using mutual information.

5. To provide a demonstration of how TVB can be used to evaluate unimodal and multimodal neuroimaging methods.

Originally, the primary aim of this thesis was to further develop 'trial-by-trial' parallel ICA (see Section 1.5.1.2) by addressing some of the limitations in the existing methods mentioned in Section 1.5.1.2.5. However, during the course of this project, it became clear that an evaluation framework was required to systematically guide the method development and, as no such framework currently existed, the focus of this project shifted to the evaluation of ICA and parallel ICA rather than just the development of parallel ICA.

As the accuracy of parallel ICA depends on the accuracies of single-modality (EEG and fMRI) ICAs, since ICs are first obtained from these decompositions and then matched across modalities in parallel ICA, **Aim 1** was set to evaluate the extent to which the single-modality ICAs provide source information.

As mentioned in Section 1.5.1.2.5, 'trial-by-trial' parallel ICA has not been evaluated using synthetic data. Also, the performance of a basic, 'stripped-down' version of 'trial-by-trial' parallel ICA has never been demonstrated even though, arguably, this is the cornerstone of the approach and should be used as the benchmark against which any further development to parallel ICA should be measured. Therefore, **Aim 2** was set to develop and evaluate this basic 'trial-by-trial' parallel ICA method.

Aim 3 is inline with the original primary aim to further develop 'trial-by-trial' parallel ICA, and was based on the fact that the matching of ICs across modalities has never been performed using spatial features and has only been performed in the temporal domain. As there is no a priori reason to assume that spatial matching cannot be useful, and as spatial matching may potentially improve performance either independently or by complementing temporal matching, this aim was set to explore this.

Aim 4 was also set to further develop 'trial-by-trial' parallel ICA. It was based on the limitation that ICs have not been matched across modalities using nonlinear measures of

association. This aim was also set from an exploratory point of view and mutual information (Shannon, 1948) was used as the nonlinear measure (see Section E.4.4 for a brief introduction to mutual information).

Over the course of this project, it was appreciated how useful TVB is for the generation of synthetic datasets for the evaluation of ICA and parallel ICA. Therefore, **Aim 5** was added to the list of aims to highlight this and the importance of method evaluation in general.

1.8 Thesis structure

In this chapter, **Chapter 1**, an introductory background is provided for the content of the rest of the thesis. In **Chapter 2**, single-modality EEG and fMRI ICA are evaluated separately in terms of their accuracy at extracting source information. This chapter is structured as a self-contained manuscript with the intention of submission for publication. It addresses Aim 1 and partially addresses Aim 5. Some further investigations were performed after the completion of this chapter and they are reported in **Appendix A**. These will not be submitted with the manuscript. In **Chapter 3**, it is highlighted that the method used for evaluating ICA in Chapter 2 did not account for source noise, and the implications of this are discussed. This work was a slight digression from the main aims of this thesis, but proved to be a useful exercise. In **Chapter 4**, the 'stripped-down' version of 'trial-by-trial' parallel ICA is evaluated and extended to match the ICs across modalities using spatial features and mutual information. This chapter is also structured as a self-contained manuscript with the intention of submission for publication. It also addresses Aim 1, along with Chapter 2, as well as Aims 2, 3, and 4, and partially addresses Aim 5. The supplementary material for this chapter is in **Appendix C** and this will also be submitted with this manuscript. **Chapter 5** contains a summary and general discussion about the work in this thesis and sums up its novel contributions.

Suggested reading pathway While the reader can go through the chapters sequentially, the author recommends reading Chapter 2 first followed by Chapter 4, as these two chapters are structured as independent manuscripts and they cover the main scientific contributions of this project. This can then be followed by Appendix A, Chapter 3 and then Chapter 5. This chapter, Chapter 1, may be skipped as the background for the other chapters is provided within them in sufficient detail. The appendices, apart from Appendix A, may also be skipped.

Chapter 2

Evaluation of ICA for Extracting Source Information From Simulated EEG and fMRI Signals

Abstract

Independent component analysis (ICA) is a method used to decompose a multivariate signal into linear combinations of maximally statistically independent components (ICs). When applied to electroencephalogram (EEG) and functional magnetic resonance imaging (fMRI) signals, the ICs are typically interpreted as providing some biological information about the sources that generated the signals. To better understand the conditions under which ICA provides accurate source information, a framework was developed in which sources were defined as networks of regions having instantaneously synchronised local field potential activity and the corresponding EEG and fMRI signals were simulated. Temporal ICA (tICA) was applied on the EEG signals and spatial ICA (sICA) on the fMRI signals, and ICA performance was measured for each modality in terms of how closely the IC time-courses matched those of the sources. This framework was used to evaluate how ICA performance varied as a function of the number of source networks, the size of a source network, and the level of source noise dispersion. The results showed that EEG tICA performance increased with the number of sources and both EEG tICA and fMRI sICA performances decreased with the level of source noise dispersion and varied without clear trends with source network size. This study highlights that ICs should be interpreted cautiously as, even under ideal conditions, i.e. when the sources are independent and there is no noise, they do not always accurately provide source network information.

2.1 Introduction

Independent component analysis (ICA) is a statistical method that is used to decompose a multivariate signal into linear combinations of maximally statistically independent components (ICs). Therefore, if the input signals are assumed to be linear mixtures of independent source signals, ICA provides a means for separating the source signals from the mixtures. ICA has been applied extensively in the analysis of electroencephalogram (EEG) and functional magnetic resonance imaging (fMRI) signals for (i) the removal of artefacts and for (ii) extracting information about the neural sources of these signals. In the latter case, there are still several

fundamental questions about how ICs should be interpreted that are unanswered, e.g. how should the biological sources that ICs can provide information about be conceptualised, and to what extent can the time-courses and spatial maps of ICs correspond to those of such sources? In this paper, in an initial attempt to answer these questions, a source was defined as a network of regions having instantaneously synchronised local field potential (LFP) activity, and a series of experiments on simulated EEG and fMRI signals were performed that used this definition to evaluate ICA performance under various conditions. The motivation of this work was to characterise the conditions under which ICs can be interpreted as providing information about such source networks.

ICA decomposes a matrix (Y), where each row of Y is a signal, into the product of a mixing matrix (M) and a source matrix (S), i.e. $Y = MS$ (Hyvärinen et al., 2004). Each row of S is an IC and each column of M contains the weights with which the ICs are linearly mixed in Y. These matrices are computed such that the ICs are maximally statistically independent from each other. There are two forms of ICA: temporal ICA (tICA) and spatial ICA (sICA). As their names suggest, tICA optimises the components to be independent in the temporal domain, whereas sICA optimises for independence in the spatial domain. The form of ICA is determined by how Y is structured: for tICA, each row in Y should correspond to a spatial location and each column to a time point, whereas for sICA the transpose of this structure should be used. Correspondingly, in tICA, each temporal IC (tIC) is a time-course and its corresponding column in M is its spatial map, whereas in sICA, each spatial IC (sIC) is a spatial map and its corresponding column in M its time-course. Typically, the form of ICA is chosen based on the dimension of the data that has more sampling points (Calhoun et al., 2001c), i.e. tICA is chosen when there are more time points than spatial points, and sICA is chosen when there are more spatial points than time points. Therefore, tICA is usually applied to EEG signals, and sICA to fMRI. While there may be cases when tICA on fMRI signals might be more appropriate, or a combination of sICA followed by tICA (Calhoun et al., 2003), as sICA has been much more frequently applied, this study only focuses on its application on fMRI signals.

ICA can provide source information when the sources are statistically independent, linearly mixed in the measurement signals, and only at most one of them has a Gaussian distribution (Beckmann, 2012; Hyvärinen et al., 2004). If EEG signals are understood to be linear mixtures of local field potential (LFP) signals, and if it is expected that some of the functional networks in the brain have instantaneously independent LFP signals, then it is reasonable to expect EEG tICs to correspond to such source networks. These sources may be dipolar (Delorme et al., 2012; Makeig et al., 1996; Makeig, S., et al., 1997) and generators of event-related potentials (Bagshaw and Warbrick, 2007; Makeig, 2002; Makeig, S., et al.,

1997). Similarly, fMRI sICs can also be expected to correspond to spatially independent functional networks, as each fMRI volume can be understood to be a linear mixture of the activations of spatially independent source networks. The sources may correspond to task-related activity (Calhoun et al., 2001a, 2008) and resting-state activity (Calhoun et al., 2008; Calhoun and Lacy, 2017; De Luca et al., 2006; Greicius et al., 2007; Risk et al., 2014).

Concurrent recording of EEG and fMRI signals was first performed around twenty-five years ago by Ives et al. (1993) and has since become a fairly common method for brain scanning (Laufs, 2012). This, potentially, can enable us to observe the activity of source networks with the high spatial resolution of fMRI and the high temporal resolution of EEG. To this end, ICA has been applied within various frameworks to combine EEG and fMRI signals (Calhoun et al., 2009; Groves et al., 2011; Lei et al., 2012). Two popular frameworks are joint ICA and parallel ICA. In joint ICA, features of the EEG and fMRI signals, such as event-related potentials and statistical parametric maps, are concatenated in a single matrix and ICA is then applied to decompose that matrix (Moosmann et al., 2008), whereas in parallel ICA, the decomposition is performed separately on each modality, either dependently (Liu and Calhoun, 2007; Vergara et al., 2014) or independently (Eichele and Calhoun, 2010; Eichele et al., 2008), and the ICs are paired across modalities based on their statistical similarities.

While these frameworks present interesting ways of combining the multimodal signals, the interpretation of the joint and parallel ICs remains unclear. The components of joint ICA contain an EEG segment and an fMRI segment concatenated together. By virtue of belonging to the same component, the EEG and fMRI segments are interpreted to be associated with each other. In parallel ICA, on the other hand, EEG and fMRI components are paired with each other when they are highly correlated. Can these components be interpreted as corresponding to the activities of source networks with the EEG segments/components describing the temporal activity and the fMRI segments/components describing the spatial activity? To answer this question, a better understanding of how well ICA provides source information within each of these modalities is needed, and comparisons with 'ground-truth' EEG-fMRI datasets for which the underlying source activations are known has to be done. While this work does not evaluate these multimodal methods, it investigates the extent to which ICA provides source information when applied on each of these modalities, and it also produces datasets which could be used for their evaluation. The biological plausibility of simulated datasets can always be questioned, as they rely on simplified models of the brain and the measurement devices. However, until human brain imaging techniques advance to the point that intracranial EEG, scalp EEG, and fMRI signals can be conveniently concurrently recorded, they are likely to be the only viable option for obtaining 'ground-truth' datasets.

Some work has been done on evaluating the performance of ICA at decomposing EEG and fMRI signals into components that correspond to underlying sources. In some of these papers, novel ICA methods are introduced (Bridwell et al., 2016; Groves et al., 2011; Moosmann et al., 2008), and in others existing methods are evaluated (Calhoun et al., 2001c; Correa et al., 2007; Erhardt et al., 2011; Ghahremani et al., 1996; Risk et al., 2014). However, the simulations performed in these studies are fairly limited in terms of their biological plausibility. For example, in Ghahremani et al. (1996), EEG signals were simulated as being generated by six audio signals that were mixed using a three-shell spherical head model, and in Moosmann et al. (2008), fMRI signals were simulated as 2D spatial maps. To obtain simulated EEG-fMRI datasets that are more biologically plausible, in this study, The Virtual Brain (Sanz-Leon et al., 2013) was used. This human brain simulator implements neural field models that account for brain connectivity at multiple temporal and spatial scales, and provides concurrent multimodal measurements of brain activity, such as LFP, EEG, and fMRI.

As described earlier, the motivation of this study is to identify the conditions under which ICs obtained from EEG and fMRI signals can be interpreted as providing information about sources that are networks of regions having instantaneously synchronised LFP activity. In this study, ICA performance was evaluated as a function of the number of source networks, the size of a source network, and the level of source noise dispersion. These parameters are described in detail in the methods section. It was computationally infeasible to run a single set of simulations that varied each of these parameters in a factorial design, as that required an enormous amount of simulations. Therefore, multiple experiments with reduced parameter sets were performed sequentially, and the results of the earlier simulations informed the parameter selection of the subsequent ones.

2.2 Materials and methods

2.2.1 The Virtual Brain (TVB)

The Virtual Brain (TVB) (Sanz-Leon et al., 2015b, 2013) was used to simulate concurrent human EEG and fMRI signals. In this section, a brief overview of the framework underlying TVB and the simulation parameters used is provided. The interested reader is directed to Sanz-Leon et al. (2015b) for more details on TVB.

TVB models the brain as a dynamical system and simulates signals by integrating activity over multiple temporal and spatial scales. The smallest scale is that of a neural mass (Deco et al., 2008), which represents a spatially localised group of neurons. The model for the

Table 2.1 Connectivity parameters. Default values provided in the simulator were used for all parameters.

Parameter	Value
Long-range coupling function	Linear
Long-range coupling strength	0.00390625
Long-range coupling conduction speed	3 mm/ms
Local connectivity kernel	Gaussian
Local connectivity amplitude	1
Local connectivity variance	1 mm
Local connectivity cut off distance	40 mm

temporal dynamics of a neural mass comprises one or more state variables, which encode the cumulative activities of the neurons and their interactions. A network of neural masses is described by a neural mass model. In TVB, the cortical surface is split into vertices, where each vertex corresponds to a cortical patch, and its temporal dynamics are described by a neural mass model. Implementations of several neural mass models are provided, such as the generic 2D oscillator model, Wilson and Cowan model (Wilson and Cowan, 1972) and the Wong and Wang model (Wong and Wang, 2006). All the simulations in this study used the generic 2D oscillator model because of its ability to produce many of the dynamics observed in neuronal populations (Sanz-Leon et al., 2015b).

TVB implements long-range connectivity and local connectivity. Long-range connectivity is between brain regions and is defined using a connectome that has two matrices: one for the coupling strengths between the regions and another for the tract lengths. The coupling between regions is not instantaneous and the tract lengths are used for calculating coupling delays. Each region comprises several hundred cortical vertices, which have local connectivity. The local connectivity is defined by an exponentially decaying function that maps the effect of the activation of each vertex to that of the others within its neighbourhood. Local connectivity, unlike long-range connectivity, is instantaneous. The connectivity parameters used in the simulations in this study are listed in Table 2.1. For all of these parameters, default values provided by the simulator were used, as it was not computationally feasible to vary the parameters and choose optimum values.

In this study, the TVB demonstration dataset (Sanz-Leon et al., 2015a,b) was used to provide the anatomical information of the brain that was simulated. In this dataset, the brain is parcellated into 76 cortical regions using an atlas for the macaque brain (CoCoMac, Bakker et al. (2012)) that was further adapted to the human brain. The coupling strengths for the long-range connections are also from CoCoMac, except the callosal connections, which

are from diffusion spectrum imaging data. They are undirected and symmetric between hemispheres. The tract lengths were primarily acquired from diffusion spectrum imaging data. In cases where data was missing, estimations based on Euclidean distance was used. Unlike the coupling strengths, the tract lengths are hemispherically asymmetric. The cortical mesh provided in the dataset divides the cortical surface into 16384 vertices.

TVB provides multiple monitors for observing the simulated signals at multiple scales. In this study, the EEG and fMRI monitors were used for simulating the signals of these modalities. The EEG monitor uses a lead field matrix to project the activity of the cortical vertices to EEG sensor space. In the simulations, the default lead field matrix provided with TVB was used. This projected the activity on to 63 scalp electrodes and 2 face electrodes (IO1 and IO2). The face electrodes were excluded from analysis as they are primarily used to measure ocular activity for artefact removal and that was not the focus of this work. The EEG signals were sampled at 1000 Hz. The BOLD monitor generates an fMRI signal for each cortical vertex by convolving its neural activity with a haemodynamic response function (HRF) (Sanz-Leon et al., 2015b). The simulations in this study used the Volterra kernel (Friston et al., 2000) to model the HRF. Therefore, there were 16384 fMRI signals, one for each cortical vertex. They were sampled at 1 Hz (Time to Repetition $TR = 1$ s).

Source noise is implemented in TVB at the cortical vertex level using stochastic integration schemes (Sanz-Leon et al., 2015b). In this study, Heun's method for stochastic integration (see Appendix C of Sanz-Leon et al. (2015b)) was used to add white Gaussian noise to the time-courses of the cortical vertex signals. While the use of additive white Gaussian noise to emulate the unexplained variance in brain signals is simplistic and unlikely to be biologically plausible, this is commonly used in simulations as a first step towards accounting for the complexity of brain activity (e.g. Deco et al., 2009; Frank et al., 1999; Groves et al., 2011; Mainen and Sejnowski, 1995). The noise processes had zero means and their dispersions (standard deviations) were defined within the simulations. All the vertices within a simulation had the same noise dispersion. The effects of sensor noise, which is at the level of the EEG and fMRI measurements, were not investigated as, while this would be interesting, it would have added another dimension to the parameter space, which, in the first instance of this investigation, was not considered necessary.

2.2.2 Source networks

ICA was evaluated in terms of the extent to which the EEG temporal ICs (tICs) and fMRI spatial ICs (sICs) matched the activities of the ground-truth sources. The sources, in physiological terms, were defined as networks of regions having instantaneously synchronised local field potential (LFP) activity. In TVB, they were constructed as sets of regions that

were stimulated by a common stimulus. Each stimulus was defined by a time-course and spatial map (see the following sections). This time-course and spatial map was used as the ground-truth activity of the source network it stimulated. Figure 2.1 provides an example of a simulation with five sources; the source time-courses and spatial maps are shown, along with the simulated EEG and fMRI signals. As described in the following sections, the sources were constructed to have non-overlapping time-courses and spatial maps so that they were largely independent and, as consequence, weakly correlated. The motivation for this was to evaluate ICA under 'best-case' conditions, i.e. the sources were highly separable from each other based on their weak statistical dependence.

2.2.2.1 Time-courses

The source networks were stimulated by pulse trains with three pulses, representing a typical event-related experimental paradigm with 'on' encoding the presentation of a stimulus and 'off' encoding its absence. The source network time-courses, therefore, were signals evoked in response to these pulse-train stimuli. In Experiments 1 to 5, the pulses were 0.2 s (Figure 2.1a), and in Experiment 6, they were 1 s. As the haemodynamic response function was modelled as having a duration of 20 s, the interval between pulse onsets (T) was set to be 20.2 s (21 s in Experiment 6) so that pulses did not have overlapping haemodynamic responses. The phase offset of each train was chosen such that their pulses would have minimal overlap by using: $\theta_n = \theta_1 + (n-1)\Delta\theta$, where θ_n was the phase offset of the n-th train, $\Delta\theta = T/N$, N was the number of sources, and $\theta_1 = 3$ s. The first pulse of the simulation occurred at 3 s, as from visual inspection of the simulated EEG and fMRI signals it was apparent that this was sufficient time for the simulator to reach a steady-state (see Sanz-Leon et al. (2015b) for a discussion on network stability). For simplicity, for each source, the stimulus time-course was used to estimate the ground-truth source network time-course for the evaluation of ICA. Therefore, from hereon in the manuscript, source network time-course refers to the stimulus time-course and not the evoked signal.

2.2.2.2 Spatial maps

For each simulation, a source-to-region mapping was constructed. This mapping was encoded in a binary matrix, with the rows indexing the sources and the columns indexing the regions. For example, Figure 2.1b shows the source-to-region mapping of a simulation with five sources, where each source comprises a single, unique region. A value of 1 (yellow) in a cell indicates that a source comprises the corresponding region, and a 0 (blue) indicates otherwise.

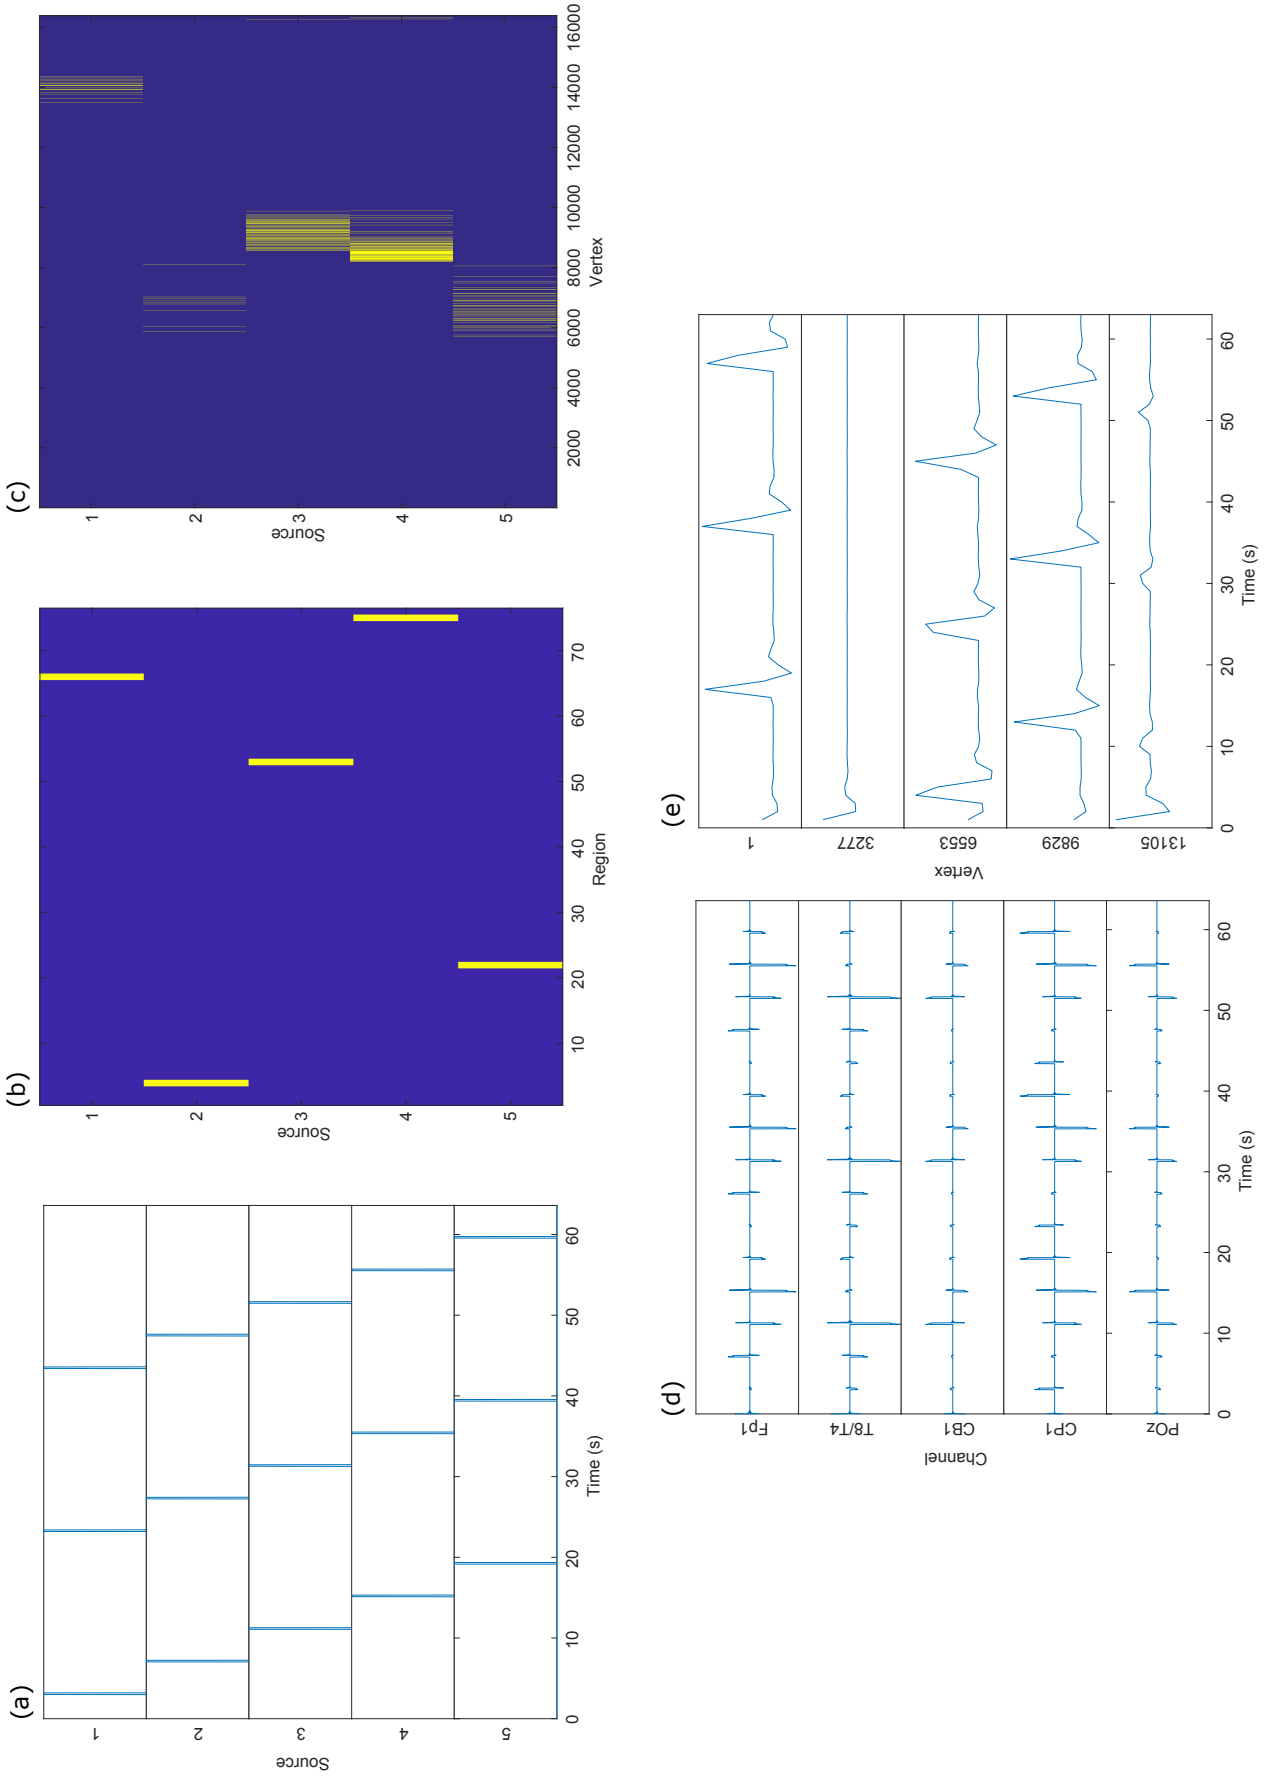
(a)
Source
Time (s)
(b)
Source
Region
(c)
Source
Vertex
(d)
Channel
Time (s)
(e)
Vertex
Time (s)

Figure 2.1 A simulation with 5 sources with no noise. (a) Source network time-courses. Each time-course is a pulse train with 3 pulses. As the number of sources is 5, the phase difference between the two consecutive pulse trains is 4.04 s. (b) Source-to-region mapping. Each source in a simulation comprised 1 to 76 regions. In this example, each source comprised a single, unique region, that was selected from a uniform distribution without replacement. A yellow cell indicates that the source comprised the corresponding region, a blue cell indicates otherwise. (c) Source network spatial maps. Each source comprised 1 to 16384 vertices. If a source comprised a region, it comprised all the vertices within that region. A yellow cell indicates that the source comprised the corresponding vertex, a blue cell indicates otherwise. (d) Simulated EEG signals for 5 of the 63 channels (sampling rate = 1000 Hz). (e) Simulated fMRI signals for 5 of the 16384 vertices (sampling rate = 1 Hz).

The spatial maps of the sources were defined at the vertex-level: if a source comprised a region, it comprised all the vertices within that region. Figure 2.1c shows the spatial maps constructed using the source-to-region mapping in Figure 2.1b. Similar to the source-to-region mapping, the spatial map of a source was encoded in a binary vector that had a value of 1 (yellow) corresponding to the vertices it comprised and a 0 (blue) for all other vertices.

Apart from the first source, which was experimentally manipulated, all the other sources comprised single, unique regions, which were selected randomly from a uniform distribution without replacement. The spatial maps of these sources were weakly correlated, as they were sparse and only contained 1s in a few non-overlapping vertices. The sources were not constructed to correspond to known functional networks in the brain. Rather, they were constructed to evaluate the performance of ICA more generally in terms of finding ICs that corresponded to networks of regions that were instantaneously synchronised irrespective of their anatomical configuration.

2.2.3 Independent component analysis (ICA)

2.2.3.1 Algorithm

Decomposing a matrix (Y) into the product of a mixing matrix (M) and a source matrix (S), i.e. $Y = MS$, such that the rows of S are mutually statistically independent is nontrivial in terms of computability, as there may not exist an S with statistically independent rows, and complexity, as obtaining the exact marginal and joint probability distributions of the rows of S is computationally expensive. Therefore, ICA algorithms estimate independence using heuristics to obtain *maximally* statistically independent components.

ICA algorithms broadly fall into two categories in terms of their estimation of independence: they either minimise mutual information, e.g. Infomax (Bell and Sejnowski,

Table 2.2 Nonlinearities $g(u)$ for FastICA

Name	$g(u)$
pow3	u^3
tanh	$\tanh(u)$
gauss	$ue^{\frac{u^2}{2}}$
skew	u^2

1995), or they maximise non-Gaussianity, e.g. FastICA (Hyvärinen and Oja, 2000). While it is clear that minimising mutual information (Shannon, 1948) would lead to maximally independent components, as mutual information is a measure of statistical dependence (see Cover and Thomas (2006) for further explanation), it has been shown that maximising the non-Gaussianity of S also provides independent components (Hyvärinen and Oja, 2000; Lee et al., 2000).

In this study, a MATLAB (Mathworks, USA) implementation of FastICA (Gävert et al., 2005) was used. Earlier versions of this implementation have been used in other studies (Vigário and Oja, 2000; Wibral et al., 2008). FastICA was selected because it is one of the most widely used and computationally efficient ICA algorithms. FastICA estimates independent components by maximising their non-Gaussianity using negentropy, which measures the difference between the Shannon's entropy (Shannon, 1948) of a variable and that of a Gaussian variable with the same variance (Hyvärinen and Oja, 2000). The FastICA implementation used provided four nonlinearities $g(u)$ for estimating the unmixing matrix (M^{-1}) based on negentropy (Table 2.2).

As components obtained simply on the basis of their non-Gaussianity may not be orthogonal, FastICA checks for the orthogonality between components during their estimation. This process can be deflationary, where components are estimated one-by-one, or symmetric, where they are estimated in parallel. In this study, symmetric orthogonalisation was used along with the tanh nonlinearity to estimate negentropy, as this parameter combination has been reported to provide optimal results (Giannakopoulos et al., 1999) and has been used in other studies (e.g. Wibral et al., 2008).

2.2.3.2 Evaluation

ICA performance in each modality was evaluated for each simulation in terms of how closely the time-courses of the ICs matched those of the source networks. The number of ICs obtained in each ICA decomposition was set to be equal to the number of sources for two reasons: (i) to facilitate the decomposition process by reducing the likelihood of

under/over-decomposition, and (ii) it was expected that for ICA to perform well there should be a one-to-one correspondence between the sources and the ICs.

2.2.3.2.1 EEG tICA performance As described in the introduction, in tICA, the tICs are time-courses, and their corresponding columns in the mixing matrix (M) are their spatial maps. Therefore, the performance of tICA was evaluated in terms of how closely the tICs matched the source network time-courses. On each set of EEG signals from each simulation, ICA was applied 30 times to account for the differences between decompositions due to variance between the initial states of the unmixing matrices. The first 2 s of the simulation were excluded so that the simulator could reach a steady-state. For each ICA decomposition, the pairwise Pearson correlation between each source time-course and each tIC was calculated (Figure 2.2) and, for each source, the correlation with the highest absolute value was retained. The magnitude, and not direction, of the correlation was used as only the amount of variance in the source time-course the IC could explain was of interest. For each source, tICA performance was defined as the mean of these correlation values across the 30 decompositions. For each simulation, the tICA performance of the first source was used as the performance measure for the entire simulation, as that was the only source that was experimentally manipulated.

2.2.3.2.2 fMRI sICA performance As described in the introduction, in sICA, the sICs are spatial maps, and their corresponding columns in the mixing matrix (M) are their time-courses. The performance of sICA was evaluated in terms of how closely the sIC time-courses matched the predicted BOLD signals of the sources. The comparison was done with the predicted BOLD signals, rather than the source network time-courses, as the sICs were obtained from fMRI signals and, therefore, their temporal resolution matched the scale of the fMRI signals and not that of the source networks.

The predicted BOLD signals were constructed by convolving the source time-courses with the same first order Volterra series haemodynamic response function (Friston et al., 2000) that was used to simulate the fMRI signals and then downsampling them to match the sampling rate of the fMRI signals. Similar to the process that was used to measure EEG tICA performance, on each set of fMRI signals from each simulation, ICA was applied 30 times. The first 2 s of the simulation were excluded so that the simulator could reach a steady-state. For each sICA decomposition, the pairwise Pearson correlation between each source predicted BOLD time-course and each sIC time-course was calculated, and then, for each source, the correlation with the highest absolute value was retained. For each source, sICA performance was defined as the mean of these correlation values across the

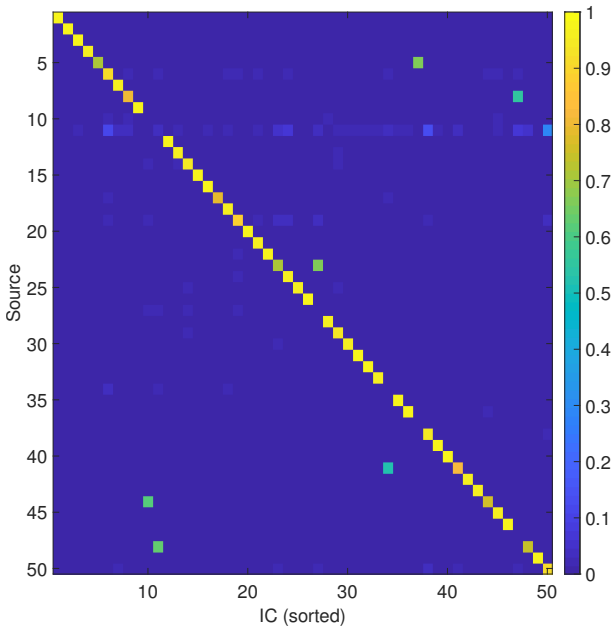

Figure 2.2 Example of pairwise correlations between source and tIC time-courses in a single ICA decomposition. Each cell contains the magnitude of the correlation between the corresponding source time-course and tIC time-course. In this simulation, there were 50 sources. (Note: for visualisation purposes, the ICs have been sorted such that the highest correlations lie on the diagonal where possible.)

30 decompositions. For each simulation, the sICA performance of the first source was used as the performance measure for the entire simulation, as that was the only source that was experimentally manipulated.

2.2.4 Experiments

Six experiments were performed. These are described in detail in the following sections. Experiment 1 evaluated ICA performance across the number of sources. Based on the results of this experiment, subsequent experiments had 50 sources. Experiment 2 evaluated ICA performance across the size of the first source network in the absence of noise. Based on the results of this experiment, the first source network had 37 regions in subsequent experiments. Experiment 3 evaluated ICA performance as a function of the source-to-region mappings. The results of this experiment did not affect the design of the subsequent ones. Experiment 4 evaluated ICA performance across source noise dispersion. Based on the results of this experiment, the source noise dispersion was set to 1 in Experiment 5, which evaluated ICA performance in the presence of source noise. Experiment 6 tested if longer pulses improved ICA performance. This had the same parameters as Experiment 5, with the exception that

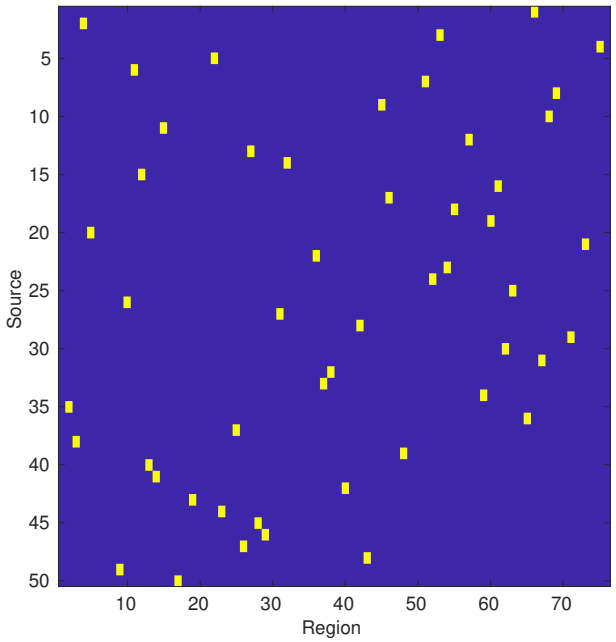

Figure 2.3 The source-to-region mappings of the simulations used in Experiment 1 to evaluate ICA performance as a function of the number of sources, which was varied from 1 to 61. A yellow cell indicates that the source comprised the corresponding region, and a blue cell indicates otherwise. Each source comprised a single, unique region. The simulation evaluating ICA performance with n sources contained the first n sources. For example, the simulation with a single source only contained the first source, the simulation with two sources contained the first two, and so on.

the pulses had a duration of 1000 ms instead of 200 ms. The data and code used in this study are available upon direct request.

2.2.4.1 Experiment 1: Evaluation of ICA performance across the number of sources

Experiment 1 investigated how ICA performance varied as a function of the number of source networks. Sixty-one simulations were run and, across them, the number of sources was varied from 1 to 61. The maximum number of sources was 61 as the simulated fMRI signals had 61 time points. Each source comprised a single, unique region, that was randomly selected from a uniform distribution without replacement (Figure 2.3). The simulations were constructed such that the simulation with n sources had the source-to-region mapping shown for the first n sources.

As described in Section 2.2.2.1, the time-courses of the sources were approximated by pulse trains with three pulses each. As the phase difference between the pulse trains was dependent on the number of sources, apart from the first source, the phase offset of each source varied across the simulations. For example, when there were 5 sources, the phase

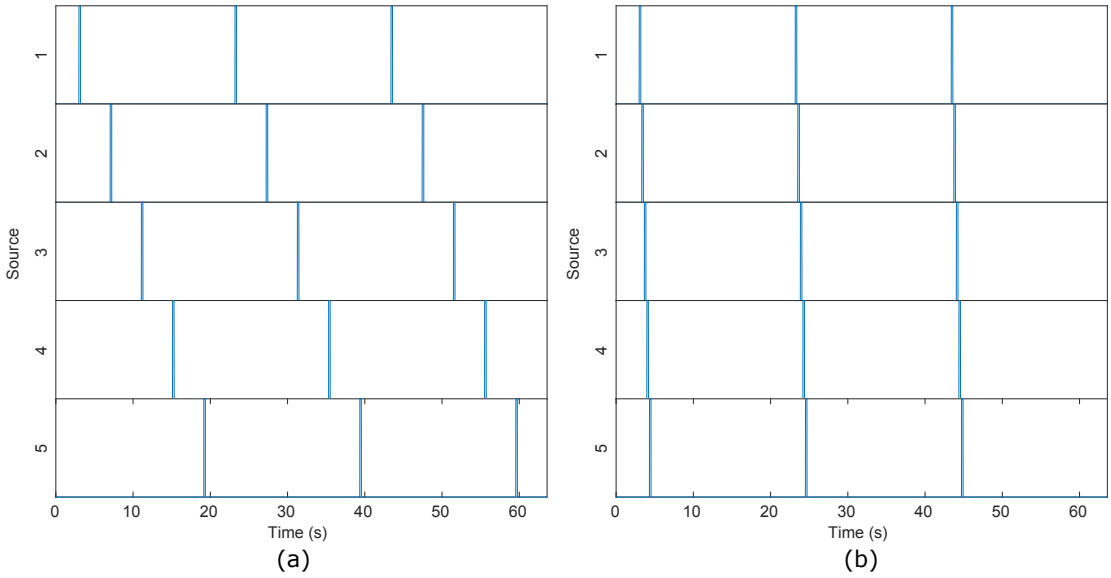

Figure 2.4 Time-courses of the first five sources in Experiment 1 used to evaluate ICA performance as a function of the number of sources in two cases: with (a) 5 sources and (b) 61 sources. The phase difference between the time-courses of two consecutive sources in (a) is 4.04 s, whereas in (b) it is 0.331 s.

difference between consecutive sources was $\Delta\theta = 20.2/5 = 4.04$ s, whereas when there were 61 sources, the phase difference was $\Delta\theta = 20.2/61 = 0.331$ s (Figure 2.4). As the performance of ICA in a simulation was always measured in terms of its performance for the first source, and the time-course of the first source remained constant across the simulations, the variation of the time-courses of the other sources was not considered to be a confound in the comparison between the simulations.

The noise dispersion was set to 0. Therefore, these simulations had parameters that were ideal for ICA: the sources were temporally and spatially independent and there was no noise.

There was not any hypothesis on how tICA and sICA performances would vary as a function of the number of sources. These simulations were performed on an exploratory basis and their purpose was to (i) see if and how tICA and sICA performances varied with the number of sources and (ii) identify the optimal number of sources to be used in the subsequent experiments.

On visual inspection of the results (Figure 2.8) it was noticed that both tICA and sICA performances tended to increase with the number of sources, and that tICA performance had a linear relationship with the number of sources. To test this linear relationship, Pearson's correlation between tICA performance and the number of sources was calculated. The

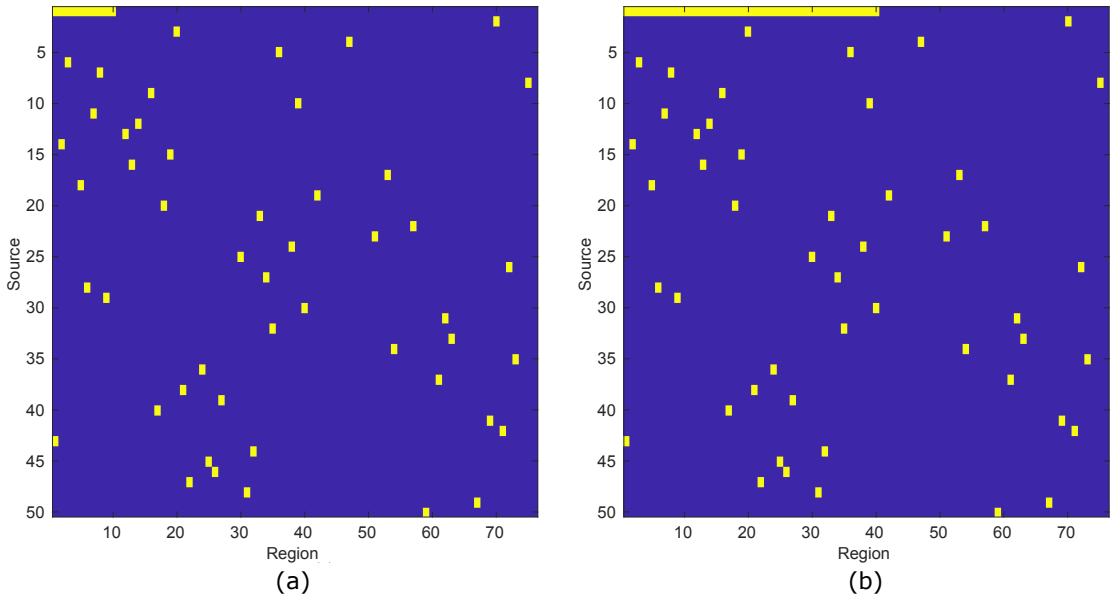

Figure 2.5 The source-to-region mappings of two of the simulations in Experiments 2, 5, and 6 used to evaluate ICA performance as a function of source network size. A yellow cell indicates that the source comprised the corresponding region, and a blue cell indicates otherwise. There were 50 sources and the size of the first source was varied from 1 to 76 regions in steps of 3. The size of the first source in (a) is 10 regions and in (b) is 40 regions. When the size of the first source was n, it included the first n regions. All the other sources comprised the same single, unique regions, across the simulations.

optimal number of sources for the simulations that followed was chosen subjectively by selecting a value for which both tICA and sICA performances were greater than 0.8.

2.2.4.2 Experiment 2: Evaluation of ICA performance across source network size in the absence of noise

A set of 26 simulations was run to investigate how tICA and sICA performances varied as a function of the size of the first source in the absence of source noise, i.e. the noise dispersion was set to 0. All the simulations had 50 sources, based on the results of Experiment 1 (Section 2.3.1). The size of the first source was varied across the simulations from 1 to 76 regions in steps of 3 regions. This was done sequentially, i.e. when the source had a size of n regions, it comprised the first n regions, as indexed in the parcellation that was used (Figure 2.5). All the other sources comprised the same single, unique regions, that were selected from a uniform distribution without replacement. The maximum size was set to 76 regions as that was the maximum number of regions provided in the parcellation.

The hypothesis was that both tICA and sICA performances would not vary as functions of source network size in the absence of noise. This was based on the understanding that the ability of ICA to separate a source into a unique component is dependent on the signal-to-noise ratio of the activity of the source in the signals that are measured. While increasing the size of a source would increase the strength of its signal, in the absence of noise, the signal-to-noise (SNR) ratio would remain constant. It could be argued that for a single source, at the sensor level, the activities of the other sources would act as noise. This was not considered to be the case as it was expected that the contribution of each source would be accounted for by each IC and, therefore, for each source, there would not be any residual noise affecting its SNR at the sensor level.

Two separate one-way ANOVAs were used to test for the main effects of size on tICA performance and sICA performance. In both cases, the conclusion that ICA performance was not affected by size was drawn only when the corresponding F-value had $p \geq 0.05$.

2.2.4.3 Experiment 3: Evaluation of ICA performance across source-to-region mappings

A set of 30 simulations was run to test if tICA and sICA performances varied for a single simulation parameter set across multiple source-to-region mappings. The motivation for this was to check if there was need to run each simulation multiple times with multiple source-to-region mappings, which would be the case if there was large variation in ICA performance across the mappings. All simulations had 50 sources, based on the results of Experiment 1 (Section 2.3.1), and the first source always comprised the first 37 regions, based on the results of Experiment 2 (Section 2.3.2). The other sources all comprised single, unique regions, that were randomly drawn from a uniform distribution without replacement. The mapping of these sources varied across the simulations. For example, the source-to-region mappings of two different simulations are shown in Fig 2.6. The noise dispersion was set to 0.

The hypothesis was that tICA and sICA performances would not significantly vary across the multiple source-to-region mappings. Two separate one-way ANOVAs were used to test for the main effects of the source-to-region mappings on tICA performance and sICA performance. In both cases, the ICA performances were defined as not being statistically significantly affected by the source-to-region mappings when the corresponding F-values had $p \geq 0.05$.

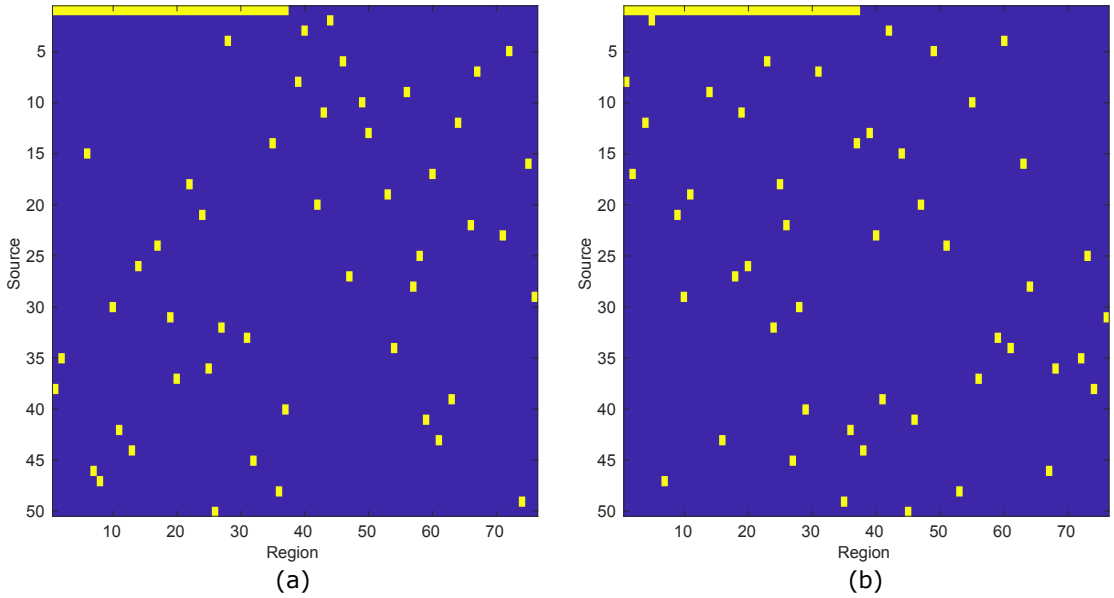

Figure 2.6 The source-to-region mappings of two of the simulations in Experiment 3 used to evaluate the robustness of ICA performance across multiple source-to-region mappings. A yellow cell indicates that the source comprised the corresponding region, and a blue cell indicates otherwise. In all of the simulations, there were 50 sources and the first source comprised the first 37 regions. All the other sources comprised single, unique regions, the mapping of which varied across the simulations. Note: in contrast, in Figure 2.5, the mapping of the first source varied across simulations and those of all the other sources remained constant.

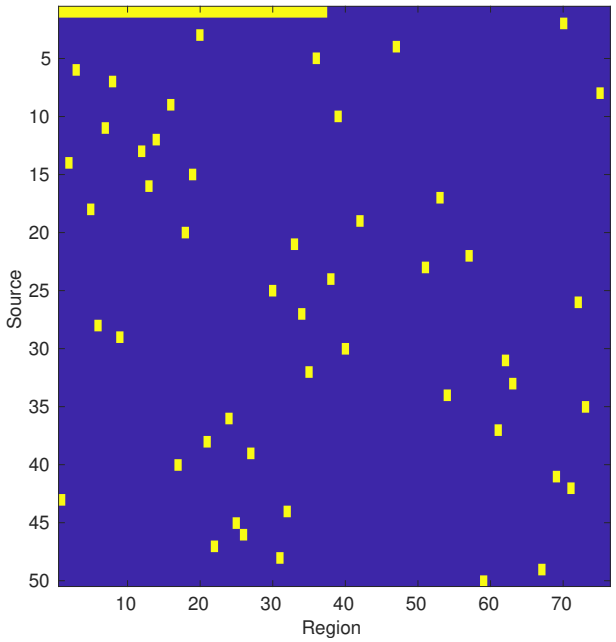

Figure 2.7 The source-to-region mapping of the simulations in Experiment 4 used to evaluate ICA performance as a function of source noise dispersion, which was varied from 0 to 2 in steps of 0.1. A yellow cell indicates that the source comprised the corresponding region, and a blue cell indicates otherwise. There were 50 sources and the first source comprised the first 37 regions. All the other sources comprised single, unique regions.

2.2.4.4 Experiment 4: Evaluation of ICA performance across source noise dispersion

A set of 21 simulations was run to investigate how tICA and sICA performances varied as a function of the the source noise dispersion, which ranged from 0 to 2 in steps of 0.1. The simulations had the same source-to-region mappings (Figure 2.7): there were 50 sources, the first source comprised the first 37 regions, and all the other sources comprised single, unique regions that were drawn from a uniform distribution without replacement. The number of sources was to set to 50 based on the results of Experiment 1 (Section 2.3.1), and the size of the first source was set to 37 regions based on the results of Experiment 2 (Section 2.3.2).

The hypothesis was that both tICA and sICA performances would decrease with the noise dispersion. The performance curves for both modalities were first visually inspected. In both cases, it was apparent that there was a linear decrease in performance when the noise dispersion was 0.1 or greater (Figure 2.11). To test this, for each modality, a linear model was fit using ordinary least squares regression, with the ICA performance for that modality set as the dependent variable, and the noise dispersion as the independent variable. Both models contained intercept terms.

2.2.4.5 Experiment 5: Evaluation of ICA performance across source network size in the presence of noise

A set of 26 simulations was run to investigate how tICA and sICA performances varied as a function of the size of the sources in the presence of source noise. The noise dispersion was set to 1. This value was selected somewhat arbitrarily; the results of Experiment 4 (Section 2.3.4) indicated that both tICA and sICA performances decreased linearly with noise dispersion and, based on this, the middle value between 0 and 2 was chosen as a representative case. The simulation parameters were the same as those used for the evaluating ICA performance in the absence of noise (Section 2.2.4.2), except that the noise dispersion for all sources was set to 1 instead of 0.

The hypothesis was that both tICA and sICA performances would increase with source network size in the presence of noise. This, again, was based on the understanding that the ability of ICA to separate a source into a unique component is dependent on the signal-to-noise ratio of the activity of the source in the signals that are measured. As increasing the size of a source would increase the strength of its signal, the signal-to-noise ratio would correspondingly increase as well, improving the ability of ICA to separate the source into a component.

The tICA and sICA performance curves were first visually inspected to see if there were apparent differences in performances across size. On visual inspection, as it appeared that there were differences (Figure 2.12a) and that there were not any clear trends, no further statistical analysis was performed.

2.2.4.6 Experiment 6: Evaluation of ICA performance across source network size in the presence of noise with 1000 ms pulses

To perform an initial investigation into if ICA performance improved when the stimulus duration was increased, Experiment 5 was repeated with the only difference being that the pulses had a duration of 1000 ms instead of 200 ms. The hypothesis was that both tICA and sICA performances would be higher for each size than they were in Experiment 5 and that they would follow the same trends as functions of size as they did in Experiment 5. ICA performances were expected to be higher because each source would have more active time periods than non-active ones compared to Experiment 5. The trends were expected to be the same as the relationship between the performances and source network size was not expected to have changed.

The tICA and sICA performances across size were compared with those in Experiment 5 to test if they were higher using a one-tailed paired t-test. Pearson correlation was calculated

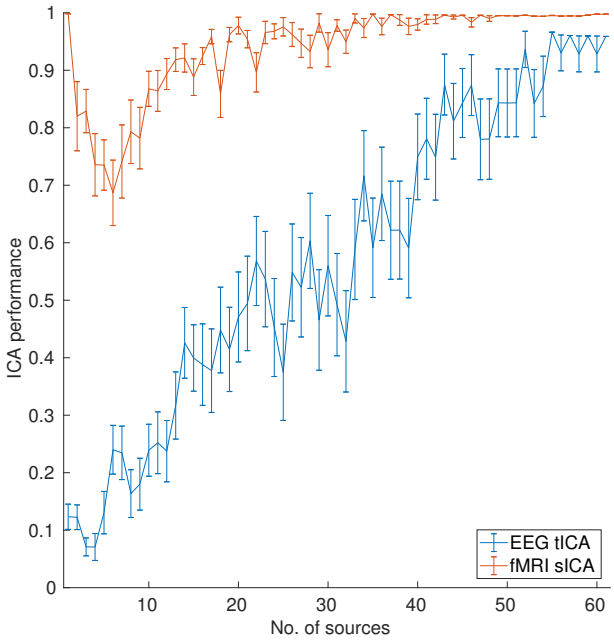

Figure 2.8 Experiment 1: ICA performance as the number of sources was varied from 1 to 61. All sources comprised single, unique regions. Each point represents the mean ICA performance across 30 decompositions, and error bars indicate standard errors of the means.

between the tICA curves and the sICA curves in both experiments to test if they had the same trends.

2.3 Results

2.3.1 Experiment 1: Evaluation of ICA performance across the number of sources

The highest tICA performance was 0.976, when the number of sources was 55, and the highest sICA performance was 0.998, when the number of sources was 1 (Figure 2.8). The performance of sICA was greater than 0.8 when the number of sources was 10 or more or 3 or less, and the performance of tICA was greater than 0.8 when the number of sources was 43 or more, with the exceptions when it was 47 and 48. The performance of tICA correlated strongly with the number of sources ($r(59) = .97, p < .001$). Based on these results, the number of sources was set to 50 for the simulations that followed.

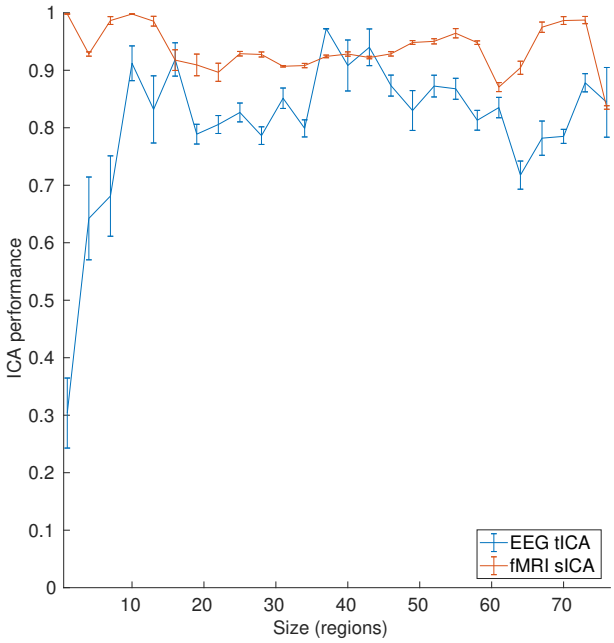

Figure 2.9 Experiment 2: ICA performance in the absence of noise across source network size. All simulations had 50 sources, the size of the first source varied from 1 to 76 regions in steps of 3, and all other sources comprised single, unique regions. Each point represents the mean ICA performance across 30 decompositions and errors bars indicate standard errors of the means.

2.3.2 Experiment 2: Evaluation of ICA performance across source network size in the absence of noise

Both tICA and sICA performances were significantly affected by the size of the first source (Figure 2.9). For tICA performance, the main effect of size was $F(25,754) = 13.29, p < .001$, and for sICA performance the main effect was $F(25,754) = 26.26, p < .001$. The performance of tICA was greater than 0.7 when the size was 10 regions or more, and sICA performance was greater than 0.8 for all sizes. The maximum tICA performance was 0.972, when the size was 37 regions, and the maximum sICA performance was 0.998, when the size was 1 region. For the simulations that followed, the size of the first source was set to 37 regions, based on the result that this size had the highest tICA performance and that the sICA performance with this size was 0.924, which was also very high.

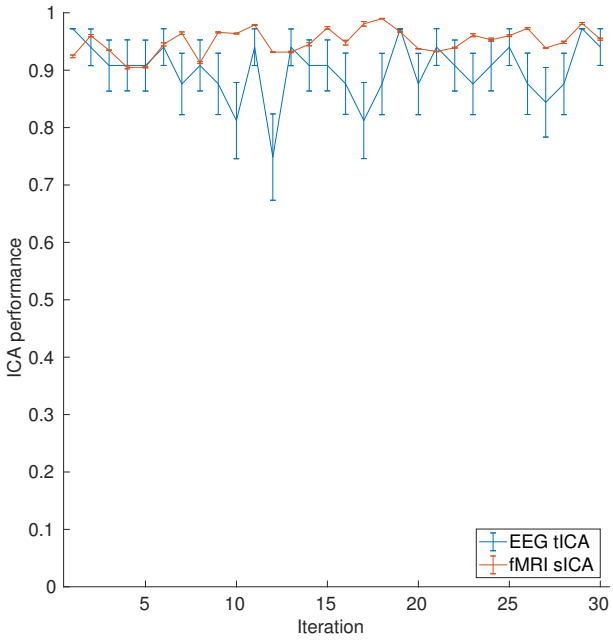

Figure 2.10 Experiment 3: ICA performance across 30 different source-to-region mappings. All simulations had 50 sources, the size of the first source was 37 regions, and all the other sources comprised single, unique regions. Each point represents the mean ICA performance across 30 decompositions and error bars indicate standard errors of the means.

2.3.3 Experiment 3: Evaluation of ICA performance across source-to-region mappings

The performance of tICA was not statistically significantly affected by the source-to-region mappings, with the main effect of the mappings having $F(29,870) = 1.19, p = .229$. However, sICA performance was statistically significantly affected by the mappings ($F(29,870) = 131.54, p < .001$). The variance of tICA performance across the mappings was 0.025 and the variance of the sICA performance was 5×10^{-4}. Despite sICA performance being statistically significantly affected by the mappings, based on these low variances and visual inspection (Figure 2.10), the conclusion was drawn that ICA performance was reasonably robust across multiple source-to-region mappings, and that this parameter did not need to be varied to evaluate ICA performance within this study.

2.3.4 Experiment 4: Evaluation of ICA performance across source noise dispersion

The performances of tICA and sICA decreased as a function of the source noise dispersion (Figure 2.11). Both tICA and sICA performances dropped sharply from when the noise

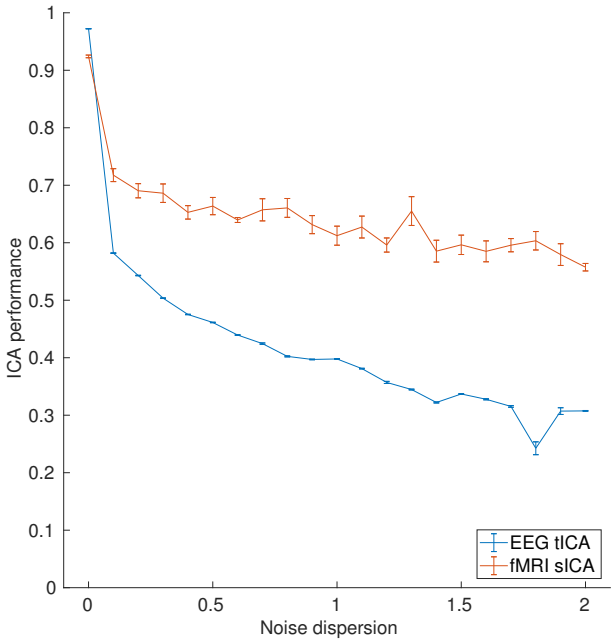

Figure 2.11 Experiment 4: ICA performance across source noise dispersion, which was varied from 0 to 2 in steps of 0.1. All simulations had 50 sources, the size of the first source was 37 regions, and all the other sources comprised single, unique regions. Each point represents the mean ICA performance across 30 decompositions and error bars indicate standard errors of the means.

dispersion was 0 to 0.1. For tICA performance, the line of best fit between the noise dispersion being 0.1 and 2 had $R^2 = .91$ and slope $\beta = -0.14$, and for sICA $R^2 = .83$ and $\beta = -0.07$. Therefore, both tICA and sICA performances decreased at fairly constant rates when the noise dispersion was 0.1 and higher, and tICA performance decreased more sharply than sICA performance. When the noise dispersion was not 0, the highest tICA performance was 0.58 when the dispersion was 0.1, and the lowest was 0.24 when the dispersion was 1.8. For sICA, the highest performance was 0.72 when the dispersion was 0.1, and the lowest was 0.56 when the dispersion was 2.

2.3.5 Experiment 5: Evaluation of ICA performance across source network size in the presence of noise

Both tICA and sICA performances were affected by source network size (Figure 2.12a). Contrary to the hypothesis, neither of the performance curves appeared to systematically increase with size. For tICA, the performance was greater than 0.39 when the size was 13 regions or greater, and was always less than 0.59. For sICA, the performance was greater

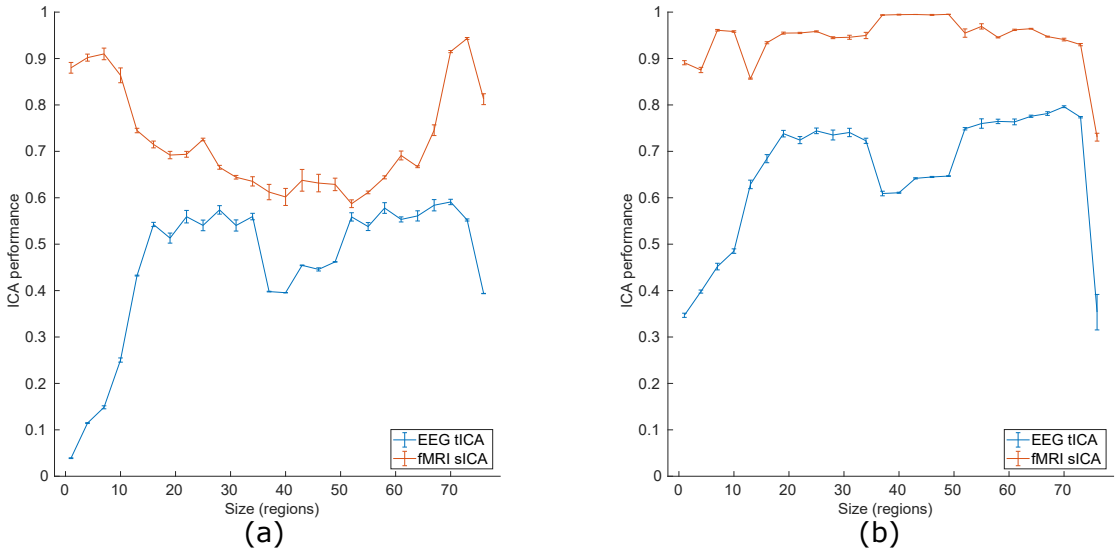

Figure 2.12 ICA performance in the presence of noise (dispersion = 1) across source network size. (a) Experiment 5: pulses had duration of 200 ms. (b) Experiment 6: pulses had duration of 1000 ms. All simulations had 50 sources, the size of the first source varied from 1 to 76 regions in steps of 3, and all other sources comprised single, unique regions. Each point represents the mean ICA performance across 30 decompositions and errors bars indicate standard errors of the means.

than 0.8 when the size was 10 regions or less, or when it was 70 regions or more, and was always greater than 0.59.

2.3.6 Experiment 6: Evaluation of ICA performance across source network size in the presence of noise with 1000 ms pulses

In the presence of noise, when the pulses were 1000 ms long, tICA and sICA performances were both higher than they had been when they were 200 ms long (Figure 2.12b). Across size, the mean increase in tICA performance was 0.1997 ($t(25) = 16.21, p < .001$) and the mean increase in sICA performance was 0.2192 ($t(25) = 7.61, p < .001$). The tICA curves had a similar trend with a correlation of 0.9138. The sICA curves did not have a similar trend with a correlation of -0.4679.

2.4 Discussion

In this study, a series of experiments were performed to characterise the performances of EEG tICA and fMRI sICA at accurately providing 'biological' source information under various

conditions, where a source was defined as a network of regions having instantaneously synchronised LFP activity. As it is difficult to obtain human EEG and fMRI datasets with 'ground-truth' source activity (e.g. intracranial EEG), the evaluation was done using simulated data. ICA performance, in both modalities, was measured using the correlations between the source time-courses and the IC time-courses. Overall, tICA and sICA showed similar trends: their performances increased with the number of sources, apart from the initial dip in sICA performance, decreased linearly with the level of source noise dispersion, did not have much variation with source network size in the absence of noise, and varied without clear trends with size in the presence of noise.

In ideal conditions, i.e. when the sources were temporally and spatially independent and there was no source noise, tICA performance improved linearly with the number of sources whereas sICA performance was consistently greater than 0.8, apart from a dip around when there were six sources. As there was not any hypothesis on how ICA performance would vary with the number of sources, it is unclear why increasing the number of sources affected ICA performance and this needs further investigation.

In the absence of noise, both tICA and sICA performances had statistically significant variation with the size of the first source network. This variation did not appear to have any systematic trends with size, except that tICA performance was very low when the source comprised a single region. Apart from this case, the variances in performance across size were low for both modalities. In the presence of noise, with the dispersion set to 1, ICA performance was measured in two cases: with the pulses being 200 ms long and 1000 ms long. While the previous experiments all used 200 ms pulses, this experiment was repeated with 1000 ms pulses to perform an initial investigation into if increasing pulse duration improved ICA performance. It had been hypothesised that both fMRI sICA and EEG tICA performances would improve when the pulse durations were longer, as the sources would have more active periods, and this was observed to be the case. The sICA and tICA performance curves had also been expected to follow the same trends. This was not the case for the sICA curve, which, as a function of source network size, had a 'u-shape' in Experiment 5 and did not show much variation in Experiment 6. The 'u-shaped' performance of sICA might be explained by the fact that the first source was the most spatially independent from the other sources when it was very small or very large - in the former case, the probability of overlap with the sources was always low, whereas in the latter case it was always high. The performance of fMRI sICA was consistently high across source network size in Experiment 6, suggesting that, with longer pulses, sICA was robust at decomposing the fMRI signals into ICs that corresponded to the sources. In the absence of noise, it had been hypothesised that ICA performance would not vary with size, whereas in the presence of noise, it was

hypothesised that ICA performance would improve with size. Apart from a small increasing trend in tICA performance in Experiments 5 and 6, the results of all the experiments that looked at ICA performance with source network size contradicted the hypotheses.

Further investigation is needed to explain these results and, perhaps, this experimental work should be done in parallel with the development of a mathematical framework that describes the relationships between ICA performance, the number of sources, and the properties of the sources. Such a theoretical framework should enable more informed interpretations of experimental results. The development of this framework, however, will be an extremely challenging task as models would have to be developed that map ICs to the signals they are extracted from that would then be combined with forward models that map source time-courses and spatial maps to EEG and fMRI signals. One of the difficulties with this would be that decomposing signals into ICs is an iterative, stochastic process and, therefore, difficult to describe by equations.

The evaluation of ICA was performed by comparing the activities of the ICs to those of the sources in terms of the stimuli that evoked responses in the signals, i.e. EEG tICs and fMRI sICs were both evaluated against stimuli time-courses. Therefore, the evaluation method measured the performance of each IC in terms of how much information it provided about stimulus-evoked responses within a source network and not how much information it provided about intrinsic activity and source noise. It is possible that a performance evaluation against the latter may have yielded better performance overall, given that stimuli time-courses and spatial maps do not account for source noise and that the ICs, to some extent, may have contained contributions from the noise processes of the sources they corresponded to. In this study, however, this evaluation was not performed as, although the noise within a source is part of its activity, it is typically not the activity of interest, especially in event-related paradigms.

The sources were constructed to be interpreted as functional networks, where all the regions within a network worked together through their instantaneous coupling. The sources had time-courses that were stimulated by non-overlapping pulse trains and, apart from the first source, which was experimentally manipulated, they all had spatial maps that comprised single, unique regions. While from a biological perspective it is implausible to have functional networks that are temporally and spatially non-overlapping, the sources were constructed this way to evaluate ICA under 'best-case' conditions, i.e. when the sources were spatially and temporally independent. As the results do not always show accurate performance of ICA under such conditions for either EEG or fMRI signals, a take-home message from this work is that one should err on the side of caution when interpreting ICs as providing source information from EEG and fMRI signals.

In terms of experimental design, the simulations could be interpreted as emulating a paradigm in which each source corresponded to a condition, and each pulse to a trial within that condition. From a real-world point of view, this would be an 'extreme' paradigm, as it would have a large number of conditions and only three trials per condition, each having a duration of 200 ms (apart from Experiment 6 in which the trials had a duration of 1000 ms). This should be considered when interpreting the results in both modalities, especially in the presence of noise. The comparison between Experiments 5 and 6 highlights that the performances would have been better if the pulse trains had longer duration and possibly also if they were repeated a greater number of times. Three pulses per train were used to keep the simulations short, around 60 s, as that was computationally feasible and deemed sufficient to provide a first look at ICA performance under the various conditions.

Choosing the number of ICs is a non-trivial problem when applying ICA (Beckmann, 2012; Calhoun and Lacy, 2017); for example, specifying more ICs than the number of sources may lead to over-decomposition (a single source split across multiple ICs), and specifying less to under-decomposition (multiple sources in a single IC). Various methods have been proposed for determining the optimum number of ICs, such as those that use information-theoretic criteria (Hui et al., 2011). In this work, this question was bypassed by setting the number of ICs equal to the number of sources as the latter were known and it was sought to evaluate ICA in ideal conditions. The framework presented in this article could be used to add to the existing body of literature on the evaluation of methods for choosing the number of ICs (e.g. Abou-Elseoud et al., 2010; Li et al., 2007) and, using the TVB pipeline, it may provide some unique insights on model order selection for specific experimental paradigms.

The results from this study suggest that, under certain conditions, both EEG tICs and fMRI sICs can be interpreted as providing information about the activities of instantaneously synchronised networks in the brain. These conditions are different for the two modalities. For sICA, the performance is generally good, and while the performance does decrease with source noise dispersion, in most cases, it still tends to be high. The same cannot be said for tICA performance, which required around 40 sources to be high, and was much lower in the presence of noise. Therefore, especially for tICA, for the ICs to correspond to sources, the paradigm should evoke activity in a large number of source networks and the level of source noise dispersion should be low. While the first parameter, to a certain degree, is within the control of the experimenter, the latter is not. Therefore, this parameter should not be manipulated, but rather set to a biologically plausible value. As, at this stage, this value is unknown, the results of this study should be interpreted as highlighting that ICA performance is more accurate when source noise dispersion is low and, therefore, ICs, especially EEG

tICs, should be interpreted as corresponding to source networks only when it can be assumed that the source noise dispersion was low.

This study also shows that the number of sources, the size of the sources, the amount of source noise, and the duration of stimulation, all affect ICA performance in both modalities. While it is intuitive why the amount of source noise and the stimulation duration would affect ICA performance, the effects of the number of sources and size were unexpected. Therefore, this study contributes to the existing body of ICA literature by identifying these effects and probing further investigation towards understanding them and their implications on experimental neuroscience. While the overall parameter space was very large, the parameter values chosen were deemed to be the most appropriate for an initial investigation of ICA performance using synthetic signals generated by TVB.

This work also makes a case for the importance of tools such as TVB for establishing stronger links between theoretical and experimental neuroscience. In the latter, there has been a radical increase in the number of methods for analysing neural signals, whereas less attention has been given to their evaluation in terms of the accuracy of the biological information they provide. This study demonstrates how these tools can be useful for performing such evaluation; even though they are limited by various simplifications of biological processes, they provide perhaps the only means of moving between the different temporal and spatial scales at which biological information is acquired, thereby being the only means forward in improving our understanding of the processes that contribute towards the signals we measure.

2.5 Conclusions

This study evaluated the performance of independent component analysis (ICA) on EEG and fMRI signals at providing information about sources that were constructed as networks of regions having instantaneously synchronised LFP activity. The motivation for this was to characterise the conditions under which ICs could be interpreted as providing information about such sources. The results showed that under ideal conditions, i.e. when the sources were spatially and temporally independent and there was no noise, spatial ICA (sICA) on fMRI signals performed well at providing source information, whereas temporal ICA (tICA) on EEG signals needed 43 or more sources to perform as well. In the presence of noise, fMRI sICA was more robust than EEG tICA and the variation in performance in both modalities with source network size did not have clear trends. These results highlight that ICs should be interpreted cautiously as even under ideal conditions there are cases when the ICs do not accurately provide source network information. While this work only focusses on ICA, it underscores the importance of evaluating methods that are commonly used in

human neuroscience to determine the extent to which the interpretation of their outputs is biologically accurate, and demonstrates how tools such as The Virtual Brain (Sanz-Leon et al., 2013) can be used to perform such evaluation.

Chapter 3

How Should ICA be Evaluated in the Presence of Source Noise?

In Chapter 2, ICA performance was measured using the correlations between the IC time-courses and the stimulus time-courses. Using this metric, ICA would only be expected to perform well if the EEG and fMRI signals could be linearly decomposed to the stimulus time-courses and spatial maps respectively. In this chapter, this assumption was first tested on the signals simulated in Chapter 2 using linear regression and it was observed that, in the absence of noise, this assumption was not violated and, therefore, ICA should have been expected to perform well in the absence of noise. However, it was also noted that, in the presence of source noise, the assumption was violated and linear decompositions of the signals in terms of the stimulus activities did not always exist. While this had been expected, as the EEG and fMRI signals were not only functions of the stimulus activities, but also of the noise, this prompted some further investigation into how ICA could be evaluated if source noise were considered to be meaningful activity (discussed below). In this chapter, the linear regression tests mentioned above are described in Section 3.1 and, in Sections 3.2 and 3.3, two alternative methods of evaluating ICA that incorporate source noise in the evaluation are explored. A summary of the chapter is presented in Section 3.4. This chapter slightly digresses from the aims of this project (see Section 1.7) and, therefore, may be skipped in the first reading of this thesis.

To recap, in Chapter 2, each source network was constructed using The Virtual Brain (TVB) (Sanz-Leon et al., 2013) as a set of regions, each of which comprised several hundred vertices (see Section 2.2.1). Within a source network, the vertices were stimulated by the same stimulus and, to the activity of each vertex, white Gaussian noise was added[1] (see

[1]The noise dispersion (see Section 2.2.1) was set to 0 in simulations in which there was no source noise.

Section 2.2.2). This was defined as *source noise* as it was at the neural level. *Sensor noise*, in contrast, would have been been added to the EEG and fMRI signals. In these experiments, sensor noise was not used, as that would have required additional simulations and that was not considered necessary in the first instance of the investigation.

Generally speaking, the sources of neural noise are not well understood (Faisal et al., 2008), and the idea behind modelling source noise is to account for the complexity that exists in brain activity in order to make the simulated signals look more realistic. While the activity is not 'meaningful', as it does not correspond to any particular functions, it is still activity that occurs at the neural level and, therefore, when extracting source information, it, arguably, could be considered relevant. That said, when analysing EEG or fMRI signals obtained from task-evoked paradigms, the task-evoked activity is typically what is of interest, and not the source noise. From this point of view, the method for evaluating ICA described in Chapter 2 is appropriate, as it measures ICA performance in terms of how much information the ICs provide about stimulus-evoked (i.e. task-evoked) activity. However, to investigate how ICA performance could be measured in terms of how much information ICs provided about the stimulus-evoked activity and the noise in the source networks they corresponded to (under the assumption that each IC is expected to match a single source uniquely), in this chapter, two methods were explored. Along with the tests with linear regression mentioned above, these are also described in this chapter.

3.1 Are the EEG and fMRI signals in Chapter 2 linear combinations of the stimulus activities?

3.1.1 Introduction

The method for evaluating ICA used in Chapter 2 (see Section 2.2.3.2) can only be expected to show ICA performing well in cases when the EEG and fMRI signals can be linearly decomposed into the stimulus time-courses and spatial maps respectively. The reasons for this are that (i) ICA performs linear unmixing and (ii) ICA was evaluated against the stimulus activities. In this section, to measure the extent to which ICA should have been expected to perform well in Chapter 2, the signals from the first five experiments[2] in the chapter were tested to see if they were linear combinations of the stimulus activities.

[2]The signals from Experiment 6 were not included in this analysis for the pragmatic reason that this chapter was completed before the experiment was added and it was decided that including the experiment would require a significant amount of time and effort and the results would be unlikely to provide further insights.

3.1.2 Methods

The objective of this investigation was to test if the EEG and fMRI signals used in Chapter 2 were linear combinations of the stimulus activities in order to determine to what extent ICA should have been expected to perform well. For a simulation, the most direct way of performing this would have been to fit a linear model to the time-course of each EEG channel in terms of the time-courses of all the stimuli and to fit a linear model to the spatial map at each fMRI time-point in terms of the spatial maps of all the stimuli. However, as in Chapter 2 ICA performance was only evaluated for the first source (see Section 2.2.3.2), a comparison between the results in the chapter with the fit of these models would not have been straightforward as they included the activities of all sources and not just the first source. Therefore, in this section, the linear models were constructed to test the extent to which the activities of the first source could be explained by linear combinations of the signals. Specifically, for EEG, for each simulation, the β-values for the model below were estimated using ordinary least squares:

$$s_T(t) = \beta_{T_1} e_1(t) + \beta_{T_2} e_2(t) + \cdots + \beta_{T_n} e_n(t) + \beta_{T_0}, \tag{3.1}$$

where $s_T(t)$ is the time-course of the stimulus corresponding to the first source, $e_i(t)$ is the time-course of the i-th EEG channel, n is the number of EEG channels, and β_{T_0} is the intercept. Similarly, for fMRI, for each simulation, the model below was estimated:

$$s_M(v) = \beta_{M_1} f_1(v) + \beta_{M_2} f_2(v) + \cdots + \beta_{M_1} f_m(v) + \beta_{M_0}, \tag{3.2}$$

where $s_M(v)$ is the spatial map of the first source, $f_i(v)$ is the fMRI volume at the i-th time-point, m is the number of fMRI volumes, and β_{M_0} is the intercept. For both modalities, linear regression performance was measured by the fit (R^2) of the model.

Even though these models only account for the stimulus activity of the first source, they test whether the EEG and fMRI signals should be assumed to be linear combinations of the stimulus activities of all sources. For illustration, consider the case of EEG: according to the proof provided in Appendix B, if the EEG signals are assumed to be linear combinations of the stimulus time-courses, then the first stimulus time-course should be a linear combination of the EEG signals. The contrapositive of this statement must also be true, i.e. if the first stimulus time-course is not a linear combination of the EEG signals, then the EEG signals should not be assumed to be linear combinations of the stimulus time-courses. The performance of linear regression tells us the extent to which the first stimulus time-course (or spatial map) is a linear combination of the EEG (or fMRI) signals and, therefore, the

extent to which the EEG (or fMRI) signals should be assumed to be linear combinations of the stimulus time-courses (or spatial maps).

3.1.3 Results

In Experiment 1 (no. of sources), linear regression performance was high ($R^2 > 0.9$) for both modalities in all the simulations (see Figure 3.1). In Experiment 2 (source network size in the absence of noise), the performances of tICA and sICA followed the same trends as that of regression in the respective modalities (see Figure 3.2). In this experiment, the performance of EEG regression was low ($R^2 = 0.46$) when the size was 1 region and high ($R^2 > 0.8$) in all other cases for both modalities. In Experiment 3 (source-to-region mappings), linear regression performance was high ($R^2 > 0.9$) for all the source-to-region mappings (see Figure 3.3). In Experiment 4 (source noise dispersion), EEG regression performance was 0.37 or lower in the presence of noise, whereas the fMRI regression performance was relatively high ($R^2 \geq 0.77$). In this experiment, the performances of tICA and sICA followed the same trends as that of regression in the respective modalities (see Figure 3.4). In Experiment 5 (source network size in the presence of noise), EEG tICA performance followed the same trend as the EEG regression performance, whereas the performance of fMRI sICA was inversely related to the fMRI regression performance (see Figure 3.5). In this experiment, EEG regression performance was always less than 0.57, and fMRI regression performance was less than 0.6 when the number of regions was 10 or less.

3.1.4 Discussion

In this section, linear models were fit to test the extent to which the stimulus activities (time-courses and spatial maps) of the first source were linear combinations of the EEG and fMRI signals. The motivation behind this was to examine the extent to which ICA should have been expected to decompose the EEG and fMRI signals in Chapter 2 into ICs such that a single IC corresponded to the activity of the first source based on the constraint that ICA assumes that the signals are linear mixtures of the sources. For this investigation, signals from the first five experiments in Chapter 2 were used and the ICA performance of each simulation was compared to the fit of the linear models mentioned in Section 3.1.2.

The performance of linear regression was mostly high in both modalities in the absence of noise (Experiments 1, 2, and 3), meaning that ICA was expected to perform well in these cases. However, in the presence of source noise (Experiments 4 and 5), the highest R^2 of EEG regression was 0.56, which was in Experiment 5 when the size of the source was 67 regions. The performance of fMRI regression was generally high and the R^2 was only less than 0.6 in

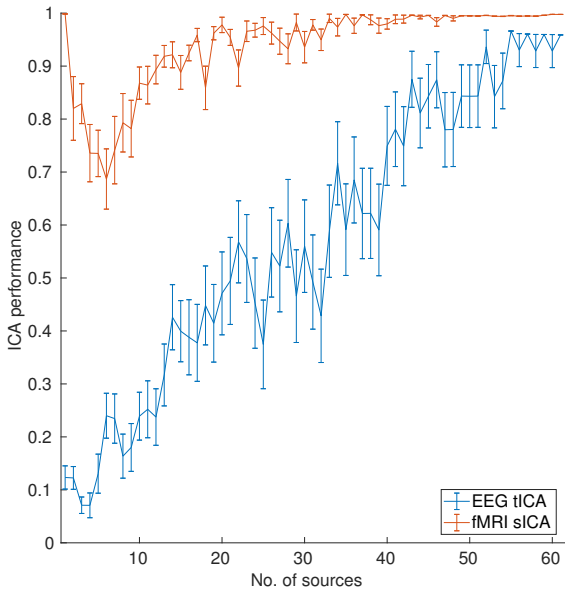

(a) ICA (across 30 decompositions)

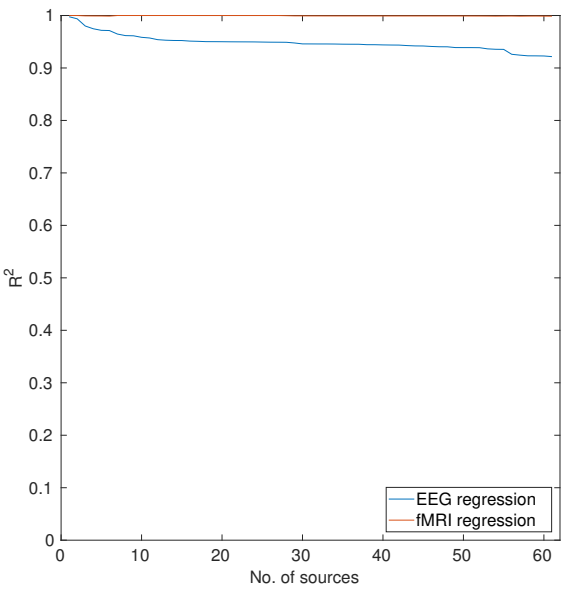

(b) Linear regression

Figure 3.1 Comparison of ICA performance of each simulation in Experiment 1 (no. of sources) in Chapter 2 with the fit of two linear models mapping the stimulus time-courses and spatial maps of the first source to the EEG and fMRI signals respectively.

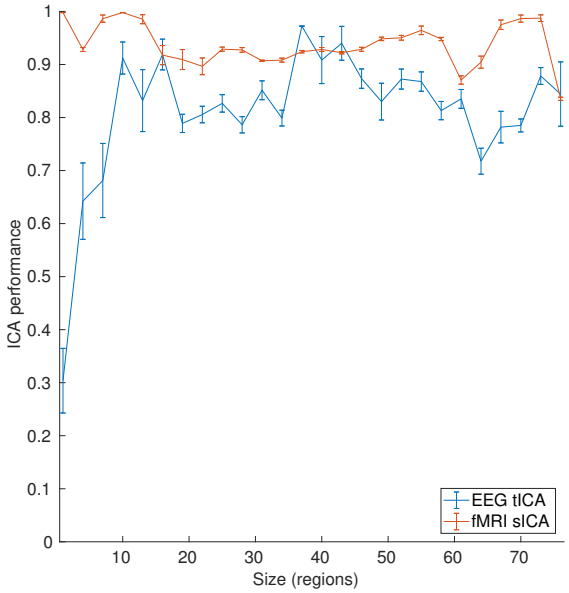

(a) ICA (across 30 decompositions)

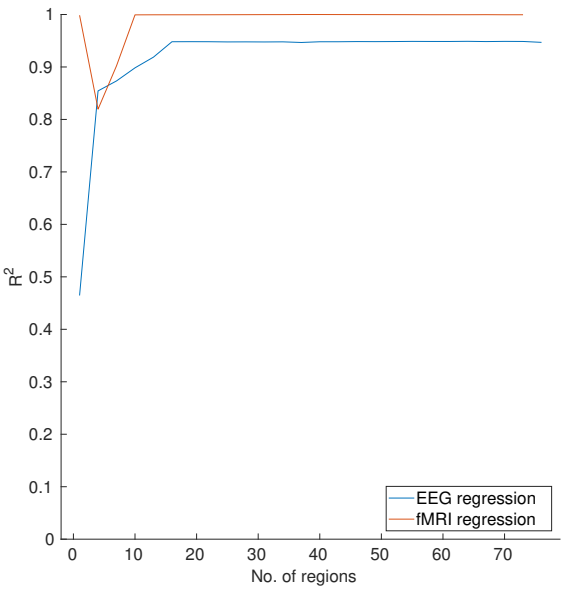

(b) Linear regression

Figure 3.2 Comparison of ICA performance of each simulation in Experiment 2 (source network size in the absence of noise) in Chapter 2 with the fit of two linear models mapping the stimulus time-courses and spatial maps of the first source to the EEG and fMRI signals respectively.

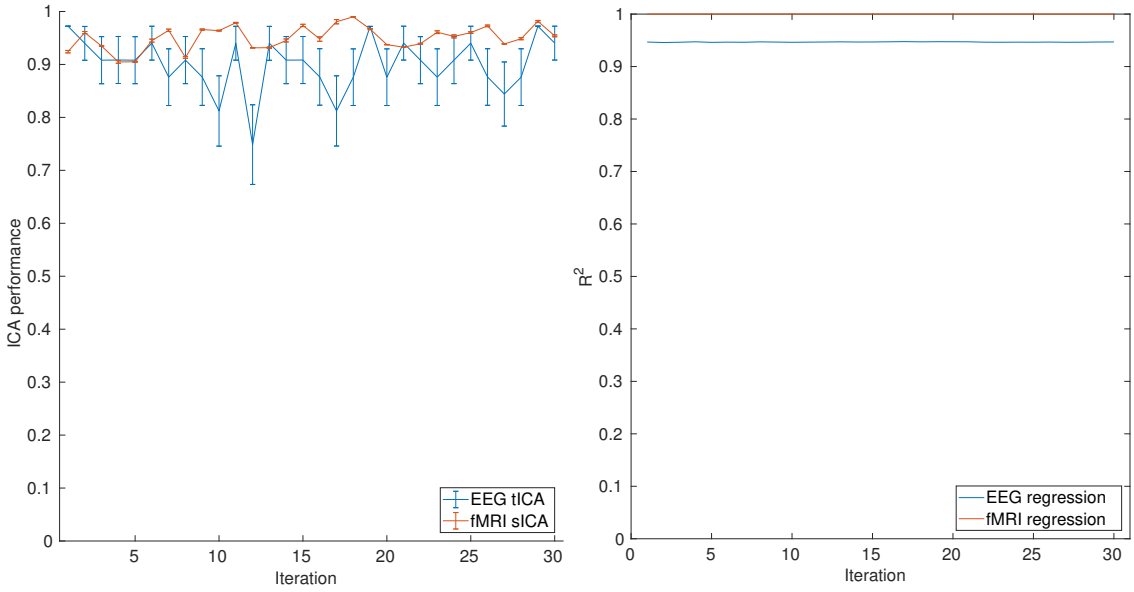

(a) ICA (across 30 decompositions) (b) Linear regression

Figure 3.3 Comparison of ICA performance of each simulation in Experiment 3 (source-to-region mappings) in Chapter 2 with the fit of two linear models mapping the stimulus time-courses and spatial maps of the first source to the EEG and fMRI signals respectively.

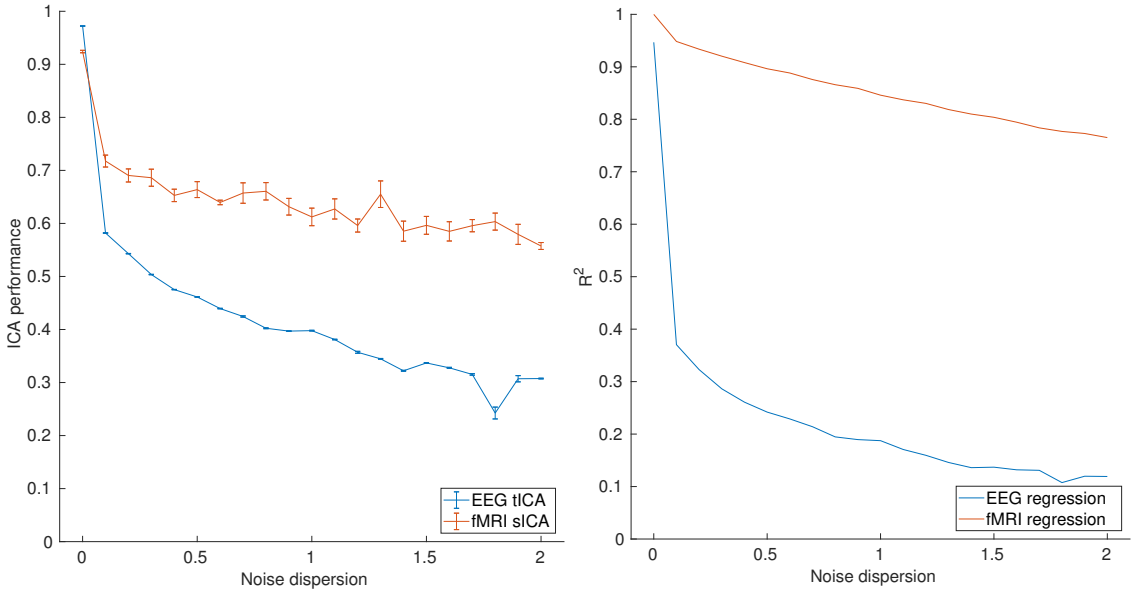

(a) ICA (across 30 decompositions) (b) Linear regression

Figure 3.4 Comparison of ICA performance of each simulation in Experiment 4 (source noise dispersion) in Chapter 2 with the fit of two linear models mapping the stimulus time-courses and spatial maps of the first source to the EEG and fMRI signals respectively.

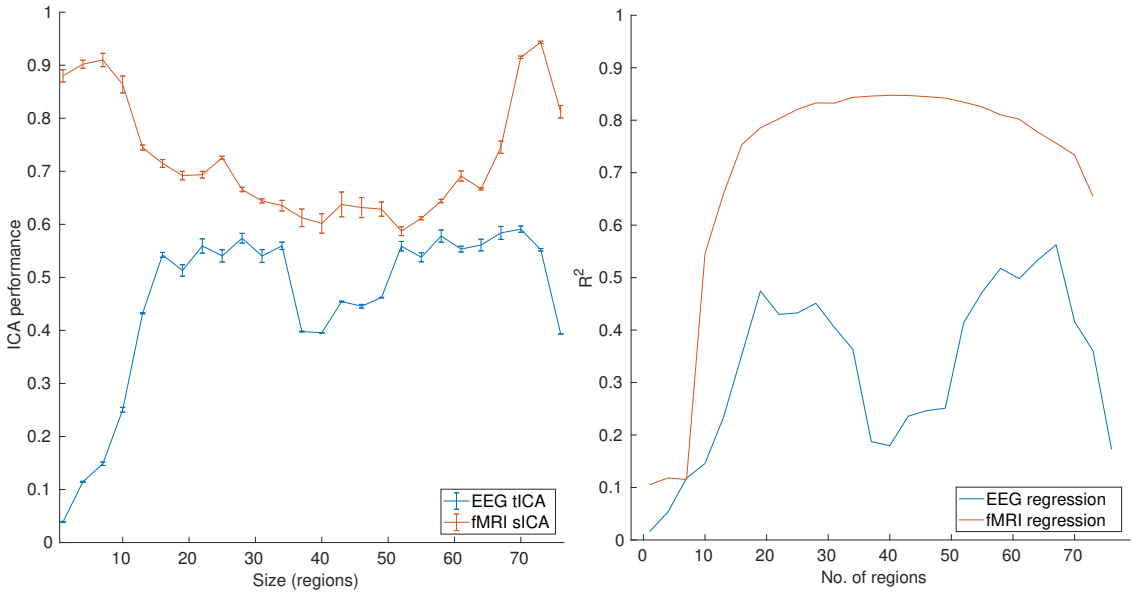

(a) ICA (across 30 decompositions)

(b) Linear regression

Figure 3.5 Comparison of ICA performance of each simulation in Experiment 5 (source network size in the presence of noise) in Chapter 2 with the fit of two linear models mapping the stimulus time-courses and spatial maps of the first source to the EEG and fMRI signals respectively.

Experiment 5 when the size of the source was 10 regions or less. In the presence of noise, the EEG ICA performance and the EEG regression performance curves were observed to have the same trends, suggesting that EEG ICA performance was strongly constrained by the fact that the signals were not just mixtures of the stimulus activities, but of the noise as well.

To elaborate further, the reason linear regression did not perform well at fitting a model to the stimulus time-course of the first source in terms of the EEG signals was because the EEG signals had variance that was accounted for by the source noise processes, whereas the stimulus time-course did not have this source of variance. Similarly, when ICA was applied on the EEG signals, some of the variance in the ICs would have been due to the source noise processes and not only because of the stimulus time-courses. By evaluating the performance of ICA in terms of the correlations of the ICs with the stimulus time-courses, the evaluation was limited to measuring the amount of source information the ICs contained about the stimulus-evoked activity and not the source noise. As, typically, in the analysis of EEG and fMRI signals only the stimulus-evoked activity is of interest, this method was appropriate for Chapter 2, which only evaluated ICA performance in this case.

The linear regression performance for the fMRI signals was high in all experiments, apart from Experiment 5. In this experiment, spatial ICA performance and fMRI linear

regression performance had u-shaped curves in opposite directions. The linear regression performance here indicated the extent to which the fMRI signals were linear combinations of the stimulus spatial maps. The opposite shapes of the curves was contrary to what had been expected, as fMRI spatial ICA performance, which was measured against the stimulus spatial maps, had been expected to follow the same trends as the fMRI regression performance. The reason for this inverse relationship is unclear. It was decided that it was not necessary to investigate this any further, as the linear regression performance curves for both EEG and fMRI had demonstrated that in the presence of noise ICA should not be expected to decompose the signals into ICs that corresponded only to stimulus activities. As, apart from this experiment, fMRI regression had performed well, further investigations in this chapter were only performed on the EEG signals and EEG ICA performance.

In Chapter 2, the performance of fMRI spatial ICA was evaluated using the stimulus time-courses and not the stimulus spatial maps. The main[3] reason for this was that the stimulus spatial maps were binary vectors and, therefore, it was not ideal to measure their similarity with the sIC spatial maps using Pearson's correlation as it assumes that both variables are continuous. Consequently, the sICs were measured for their correspondence with the sources in terms of their correlation with the predicted fMRI signals of the sources by convolving the stimulus time-courses with the haemodynamic response function (see Section 2.2.3.2.2). Despite the fact that the fMRI sICA evaluation in Chapter 2 was performed using the stimulus time-courses, fMRI linear regression performance in the current chapter was measured in terms of the extent to which the spatial map, not the time-course, of the stimulus of the first source was a linear combination of the fMRI signals. The reason for this was that spatial ICA was performed on the fMRI signals and, therefore, ICA decomposed the signals into components that corresponded to spatial maps and not time-courses.

3.1.5 Conclusions

In this section, linear regression was used to test the extent to which ICA should have been expected to decompose the EEG and fMRI signals in Chapter 2 into stimulus time-courses and spatial maps respectively. It was shown that in the absence of noise, the EEG and fMRI signals were linear combinations of *only* the stimulus activities in most cases and, therefore, ICA should have been expected to perform well, whereas in the presence of noise they were not linear combinations of *only* the stimulus activities in any case and, therefore, ICA should not have been expected to perform well. While the latter conclusion is somewhat trivial, as it was known that the signals were functions of the source noise in addition to

[3] See the last paragraph of Section 5.3 for a description of all the reasons.

the stimulus activities, it is interesting to see that, in the presence of noise, the EEG ICA regression performance curves followed the same trends as the EEG ICA performance curves, suggesting that, at least in the case of EEG, the addition of noise may have affected ICA performance only because it was not accounted for in the evaluation. In Chapter 2, it was not evaluated how much information ICs contained about the noise in the sources they corresponded to, as the focus was only on task-evoked activations. However, if source noise were considered to be of interest, as it does pertain to source activity, then the ICA performance measure should account for how much source noise information is provided by ICs. The following two sections describe two methods that were explored to account for this information. These sections only focus on the EEG signals and, correspondingly, on the EEG ICA evaluation as, apart from Experiment 5, fMRI regression performed well.

3.2 Are the EEG signals in Chapter 2 linear combinations of the mean LFP signals?

3.2.1 Introduction

In Chapter 2, the 'ground-truth' source time-courses that were used to evaluate ICA had been approximated as the stimulus time-courses, which is what are typically used to identify task-related ICs. As discussed in the previous section, this method of evaluation ignored any information the ICs provided about the 'noise' in the sources they corresponded to, which, arguably, could be considered to be relevant activity, as it pertains to the sources. In this section, to account for this, the source time-courses were approximated by the means of the LFP signals of the vertices the sources comprised, instead of the stimulus time-courses, and the same analysis with linear regression that was performed in the previous section was performed again for the EEG signals with these source time-courses. The purpose of this investigation was to assess if, in the presence of noise, the EEG signals were better explained as linear combinations of the mean LFP signals than of the stimuli time-courses because, if they were, then it would be more appropriate to evaluate EEG ICA using the mean LFP signals as the source time-courses rather than the stimuli time-courses if source noise were of interest.

3.2.2 Methods

Similar to the method described in Section 3.1.2, the models were estimated only for the activities of the first source as ICA performance was only measured for this source in Chapter

2. For the first source, the time-course was constructed as the mean LFP signal across the vertices it comprised (i.e. the vertices that were stimulated by its stimulus) and linear models were estimated to test the extent to which the activities of the source could be explained by linear combinations of the signals. Specifically, for each simulation, for EEG, the β-values of the model described in Equation 3.1 were estimated using ordinary least squares, with the only difference being that the dependent variable ($S_T(t)$) was the mean LFP signal across the vertices the first source comprised instead of the stimulus time-course. As only the source time-courses were modified and not the spatial maps, the fMRI regression was not performed again. However, for the sake of completeness, the figures in Section 3.2.3 show the fMRI regression performance curves from Section 3.1.3.

3.2.3 Results

On the signals from Experiments 1, 2, and 3 in Chapter 2, EEG regression performance was higher, and almost perfect ($R^2 \approx 1$), when using mean LFP time-courses than when using the stimulus time-courses to approximate the source time-courses (see Figures 3.6, 3.7, and 3.8). On the signals from Experiments 4 and 5, based on visual inspection, EEG regression performance was slightly higher when using mean LFP time-courses than when using the stimulus time-courses and the trends appeared to be exactly the same (see Figures 3.9 and 3.10).

3.2.4 Discussion

In this section, a comparison was performed between the fit of the linear models that explained the source time-courses in terms of the EEG signals. The models obtained in Section 3.1, in which the source time-courses were approximated by the stimulus time-courses, were compared to the models described in this section, in which the source time-courses were approximated by the mean LFP signals across the vertices the sources comprised. As mentioned in Section 3.2.1, the purpose of this investigation was to assess if, in the presence of noise, the EEG signals were better explained as linear combinations of the mean LFP signals than of the stimuli time-courses in order to determine if it would more appropriate to evaluate EEG ICA using the mean LFP signals rather than the stimuli time-courses if source noise were of interest. There were two key results that emerged from the comparisons:

- The models with source time-courses approximated by the mean LFP signals had better fits than those with them approximated by the stimulus time-courses.
- In the presence of noise, there was marginal difference between the fits of the models.

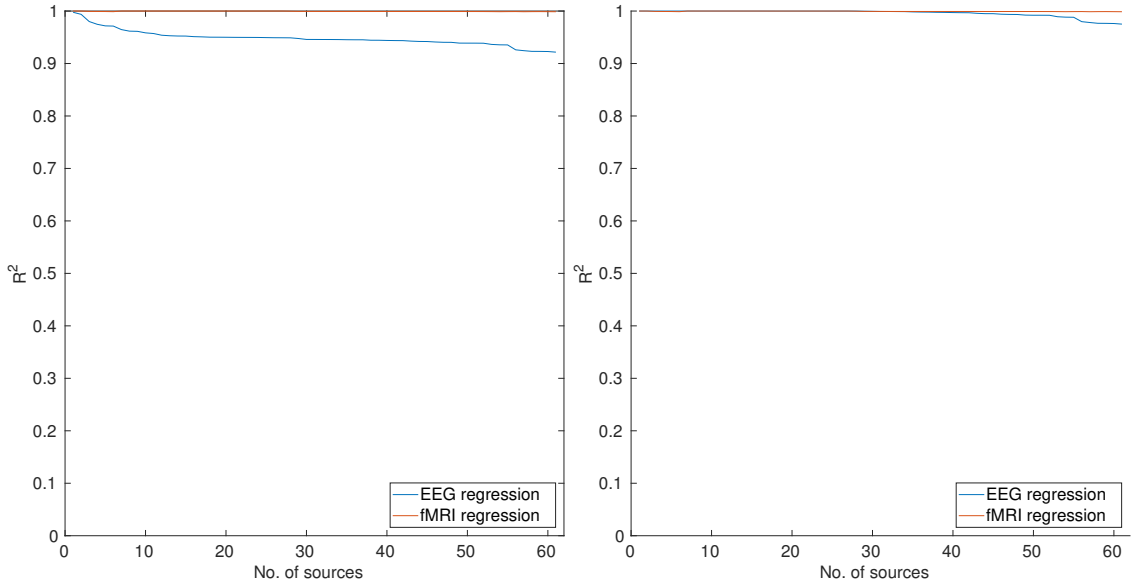

(a) Using stimulus time-courses (b) Using mean LFP signals

Figure 3.6 Comparison of linear regression performance of each simulation in Experiment 1 (no. of sources) in Chapter 2 when approximating the source time-courses by the stimulus time-courses and the mean LFP signals.

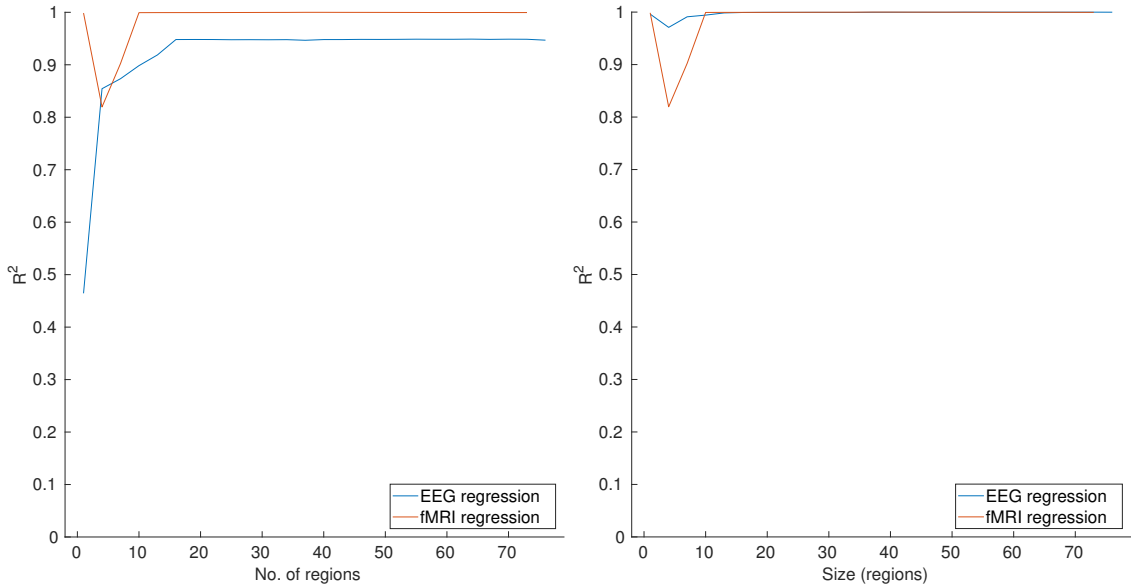

(a) Using stimulus time-courses (b) Using mean LFP signals

Figure 3.7 Comparison of linear regression performance of each simulation in Experiment 2 (source network size in the absence of noise) in Chapter 2 when approximating the source time-courses by the stimulus time-courses and the mean LFP signals.

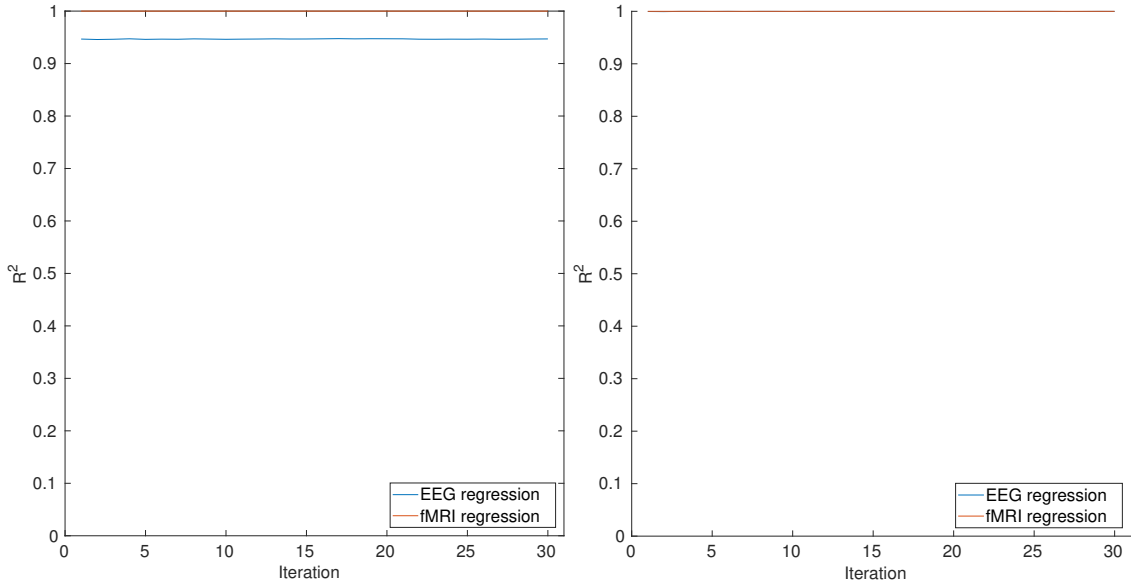

(a) Using stimulus time-courses (b) Using mean LFP signals

Figure 3.8 Comparison of linear regression performance of each simulation in Experiment 3 (source-to-region mappings) in Chapter 2 when approximating the source time-courses by the stimulus time-courses and the mean LFP signals.

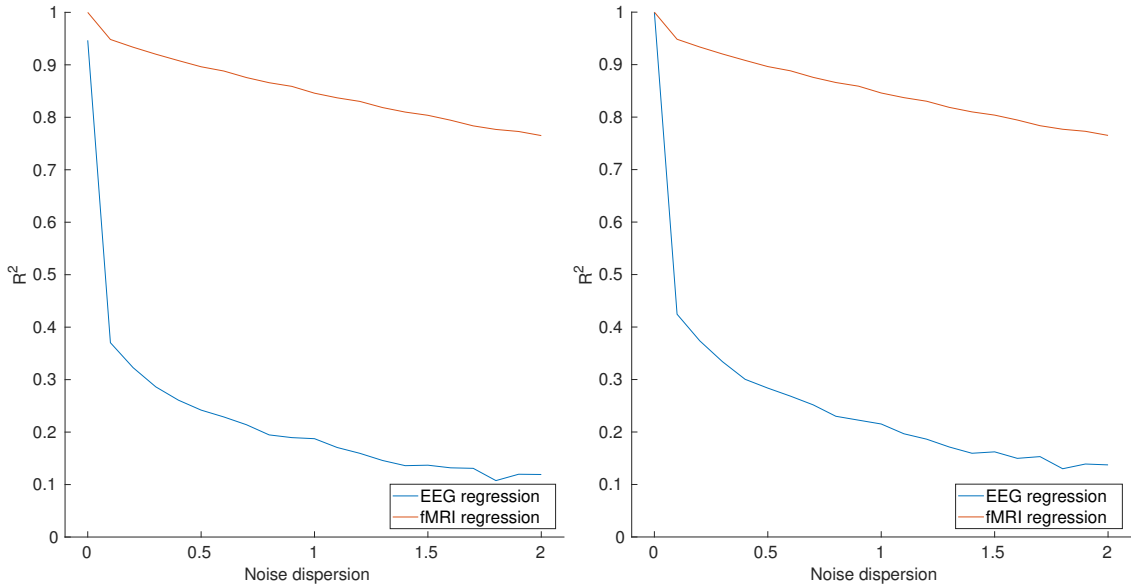

(a) Using stimulus time-courses (b) Using mean LFP signals

Figure 3.9 Comparison of linear regression performance of each simulation in Experiment 4 (source noise dispersion) in Chapter 2 when approximating the source time-courses by the stimulus time-courses and the mean LFP signals.

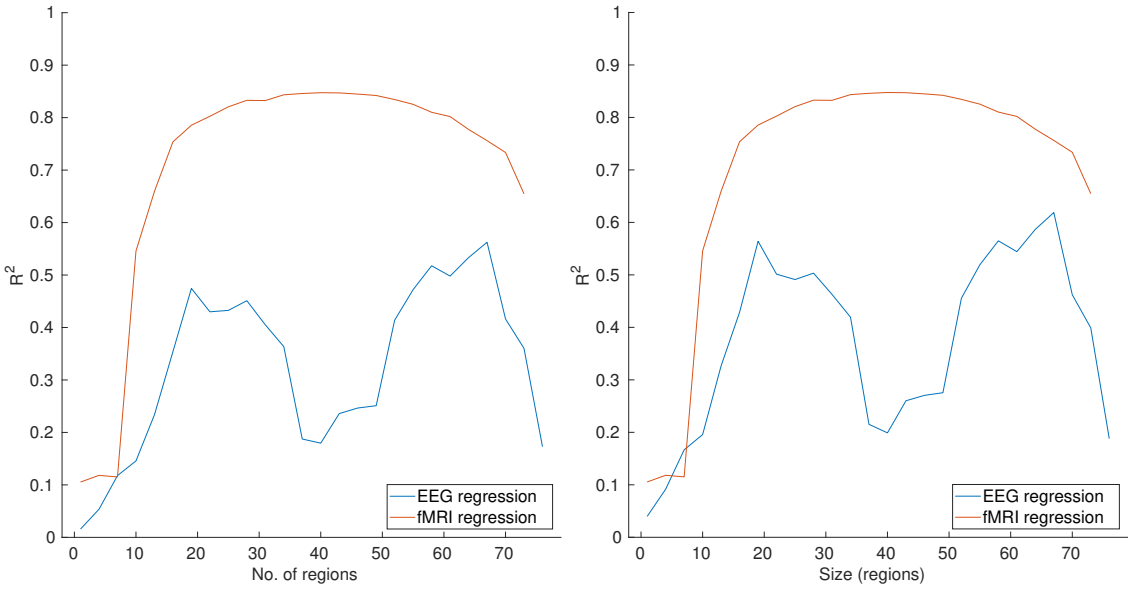

(a) Using stimulus time-courses (b) Using mean LFP signals

Figure 3.10 Comparison of linear regression performance of each simulation in Experiment 5 (source network size in the presence of noise) in Chapter 2 when approximating the source time-courses by the stimulus time-courses and the mean LFP signals.

The fact that the models with the source time-courses approximated as the mean LFP signals had better fits suggests that including the source noise in the source time-course explained greater variance in the EEG signals than the stimulus time-courses alone and, therefore, the mean LFP approximations of the source time-courses were more accurate than the stimulus time-courses. The improvement in fit could also have been due to the fact that the mean LFP signals contained the stimulus-evoked responses[4]. rather than the stimulus time-courses. However, as in the presence of noise the improvement in model fit was marginal, using the mean LFP signals as the source time-courses for ICA evaluation was not expected to provide significantly different results.

To investigate why the mean LFP signals did not result in substantially better model fits, they were visually inspected. The stimulus time-course and the mean LFP signals of the first source in Experiment 4 in Chapter 2 in simulations with the noise dispersion set as 0.1 and 1.0 are shown in Figure 3.11. Apart from the noise in the mean LFP signals, there were evoked potentials from other stimuli. These would have been generated in response to the stimuli for the 37[5] other sources the first source spatially overlapped with. However, while

[4]The difference between stimulus-evoked responses and stimulus time-courses is described in the next section (Section 3.3.2.1.2)

[5]The overlap could be with less than 37 sources, as each source, apart from the first one, comprised a single region selected without replacement from a uniform distribution of 76 regions.

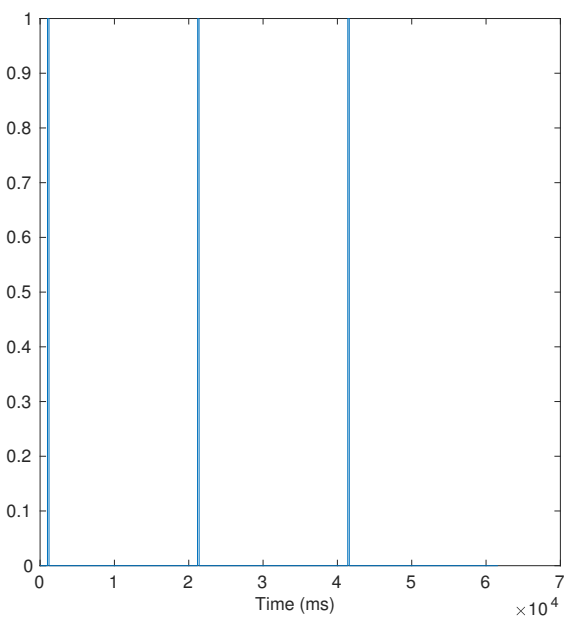

(a) Stimulus time-course

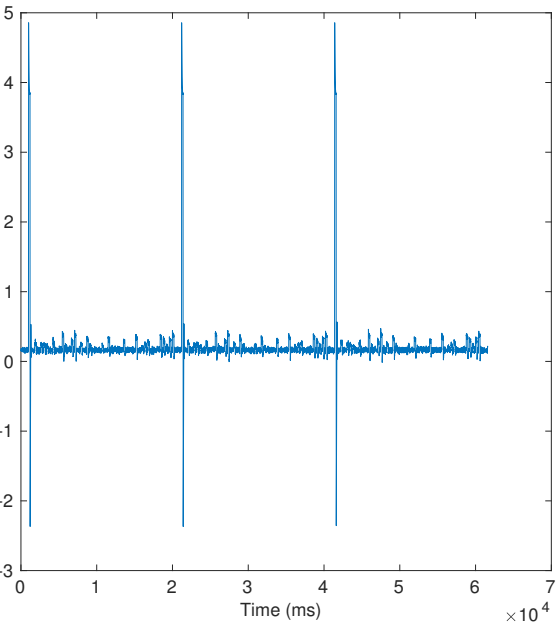

(b) Mean LFP signal (noise dispersion = 0.1)

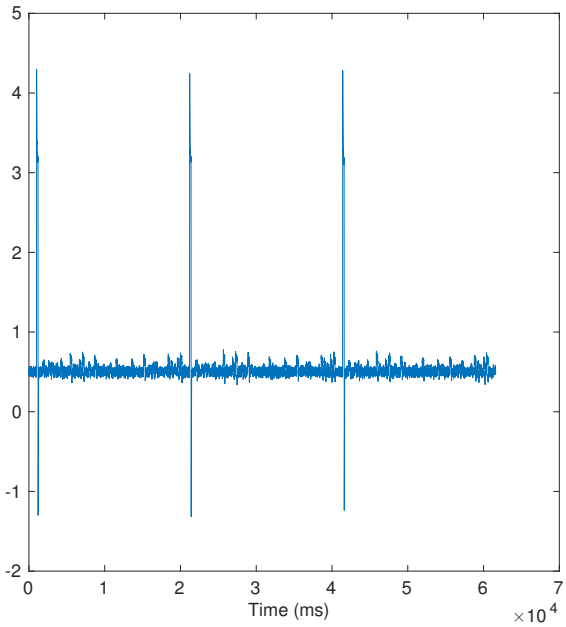

(c) Mean LFP signal (noise dispersion = 1.0)

Figure 3.11 Comparison the time-courses of the first source approximated by the stimulus time-course and the mean LFP signal.

these evoked responses should not have been part of the time-course of the first source, their presence was not expected to significantly reduce the fit of the models due to their relatively small amplitude.

One reason the models could have had a poor fit in the presence of noise was that the source noise was not completely captured by the mean LFP signal. For example, let a source network have k vertices, $v_i(t)$ be the LFP signal of the i-th vertex belonging to the source, $x(t)$ be the evoked activity in response to the stimulus, and $n_i(t)$ be the source noise added to the i-th vertex:

$$v_i(t) = x(t) + n_i(t). \tag{3.3}$$

Let $m(t)$ be the mean LFP signal of the source:

$$\begin{aligned} m(t) &= \frac{[v_1(t) + v_2(t) + \cdots + v_k(t)]}{k} \\ &= x(t) + \frac{[n_1(t) + n_2(t) + \cdots + n_k(t)]}{k}. \end{aligned} \tag{3.4}$$

The approximation of the source time-course using $m(t)$ assumes that the vertices contributed equally to the noise. This assumption, however, is incorrect, as the vertices have spatially distinct locations and, therefore, their contributions to the EEG signals would have not been the same. In fact, each EEG channel would have had a different 'ground-truth' time-course for the same source as the mixture of the noise for each vertex, for each channel, would not necessarily have been the same. To address this, in the next section, a method for evaluating ICA was explored which accounted for the individual source noise contributions of the vertices.

3.2.5 Conclusions

In this section, it was shown that using the mean LFP signal across the vertices belonging to a source, instead of the stimulus time-course, to approximate the source time-course did not provide significant improvements in fitting models to the time-course of the first source in terms of the EEG signals. Therefore, the mean LFP signal is not a significantly better approximation of the source time-course than the stimulus time-course in terms of capturing the dynamics of source noise. As discussed in Section 3.2.4, one reason for this could be that the noise of the vertices within a source are not projected exactly the same way to the EEG sensors as they have spatially distinct locations, which is an assumption made when using the mean of the LFP signals across the vertices to approximate the source time-course. Based on this, an alternative approach to evaluating ICA was investigated which evaluated

ICs in terms of the individual source noise contributions of the vertices. This is described in the next section.

3.3 Evaluation of ICA using predicted IC time-courses

3.3.1 Introduction

This section describes the results and limitations of evaluating ICA using a method that uses predicted IC time-courses. The basic idea behind this method was to evaluate each IC in terms of how much information it provided about a source in terms of the responses evoked to its stimulus as well as the noise of the vertices that belonged to it. This was applied to the EEG signals simulated in Experiment 4 in Chapter 2 with a relatively high level of noise (dispersion = 1). ICA performance was only calculated for the first ICA decomposition (out of 30) in order to get preliminary results. While the results did show some promise, it was decided not to run this method for ICA evaluation on all of the previously generated signals as it was deemed computationally infeasible (see Section 3.3.4.4).

3.3.2 Methods

This method comprised two stages. In the first stage, source noise was extracted for each vertex. In the second stage, to evaluate how closely an IC matched a source, a predicted IC time-course for each IC and source pair was constructed using the stimulus time-course and the source noise time-courses for all the vertices belonging to that source. The correlation between the predicted IC time-course and the IC time-course was used to measure ICA performance. These procedures are described in more detail in the following sections.

3.3.2.1 Source noise extraction

Two methods were used for source noise extraction. Originally, the plan had been to only use the first method. However, after performing some diagnostics on the extracted source noise, it was clear that the extraction had not been accurate, and so the second method was used.

3.3.2.1.1 Method 1: Using stimulus time-courses To extract the source noise at each vertex, linear regression was used to estimate the β-values of the following model to the LFP signal of each vertex:

$$v_i(t) = \beta_{i_0} + \beta_{i_1} s_1(t) + \beta_{i_2} s_2(t) + \cdots + \beta_{i_n} s_n(t) + \varepsilon_i(t), \tag{3.5}$$

where $v_i(t)$ was the LFP signal of the i-th vertex and $s_k(t)$ was the time-course of the k-th stimulus. The residuals of this model, $\varepsilon_i(t)$, were defined as the source noise for vertex i. The basic idea behind this model was that all the stimulus-evoked activity occurring at a vertex could be explained by the stimulus time-courses, and all other 'unexplained' activity would be source noise.

The source noise time-courses for each vertex extracted using this method are shown in Figure 3.12 for the case when there was no source noise and when the noise dispersion was set to 1. For both of these cases, it is evident that the extraction was not very accurate. In the no noise case (Figure 3.12a), the source noise should have been zero. However, it is clear that there were time-points that were non-zero. In the case with the noise dispersion set to 1 (Figure 3.12b), some systematic patterns are visible, which was not expected, as each noise observation was drawn from an independent distribution. The stimulation of the first source (which comprised 37 regions, whereas the other sources comprised single regions) is evident in both of the figures. This indicates that the stimulus-evoked responses were not completely accounted for by the stimulus time-courses used to fit Equation 3.5.

To further analyse the source noise extraction in the simulation with the noise dispersion set to 1, three vertices were randomly chosen from a uniform distribution without replacement and their histograms were visually inspected (Figure 3.13). These were noted to be positively skewed and, therefore, not normally distributed, providing further evidence that the extraction had not been accurate, as instances of the source noise were observations of a Gaussian process ($\mu = 0, \sigma^2 = 1$). This was further confirmed by one-sample Kolmogorov-Smirnov tests, which, for each vertex, were all statistically significantly different from the standard normal distribution at $\alpha = 0.05$.

One reason that the source noise extraction was inaccurate could have been that the regressors in Equation 3.5 were stimulus time-courses and not stimulus-evoked responses. This was investigated in the following section.

3.3.2.1.2 Method 2: Using stimulus-evoked response time-courses The regressors in Equation 3.5 were stimulus time-courses and not stimulus-evoked response time-courses. The difference between the two for a single pulse is shown in Figure 3.14. The reason stimulus time-courses had been used in the previous method for source noise extraction instead of the evoked responses was because they were easier to obtain than the evoked responses and were considered to be close approximations of them. To investigate if the source noise extraction had been inaccurate because the evoked responses had not been used to fit the model described in Equation 3.5, the model was fitted using the evoked responses instead. This process and the extracted source noise are described in this section.

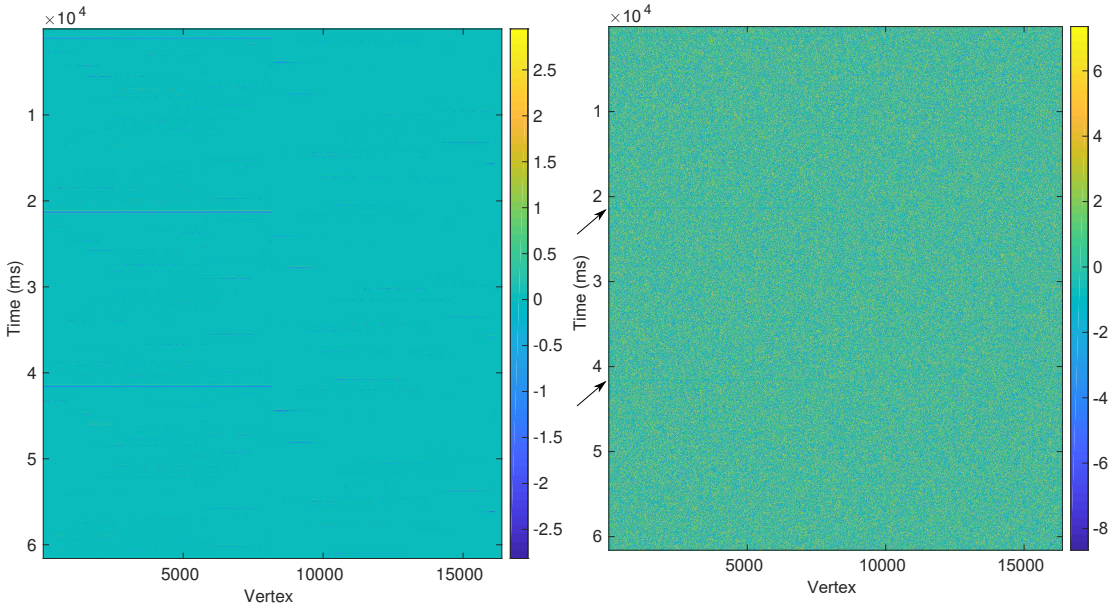

(a) Noise dispersion = 0 (no noise)

(b) Noise dispersion = 1. The arrows indicate rows with visible patterns.

Figure 3.12 Source noise extracted using the stimulus time-courses (method 1, see Section 3.3.2.1.1)

To extract the stimulus-evoked response time-course for each source, a simulation was run that only had activations in response to the stimulus that corresponded to that source. The stimulus-evoked response time-course was constructed as the mean LFP across all the vertices belonging to the source.

The source noise was extracted the same way as it had been in the previous method. The only difference between the two methods was that in the previous one the regressors were the stimulus time-courses, whereas in this one the regressors were the stimulus-evoked response time-courses. The source noise time-courses for each vertex extracted using this method are shown in Figure 3.15 for the case when there was no source noise and when the noise dispersion was set to 1. In the no noise case (Figure 3.15a), the source noise, on visual inspection, appeared to be mostly zero. However, the histograms (Figure 3.16) indicated that there were non-zero values. In the case with the noise dispersion set to 1 (Figure 3.15b), systematic patterns were not visible. However, the histograms (Figure 3.17) showed that there was still a positive skew in the distributions and that they were not Gaussian. While the source noise extraction was still not completely accurate, it was considered to be adequate enough to inspect the method for evaluating ICA using predicted IC time-courses.

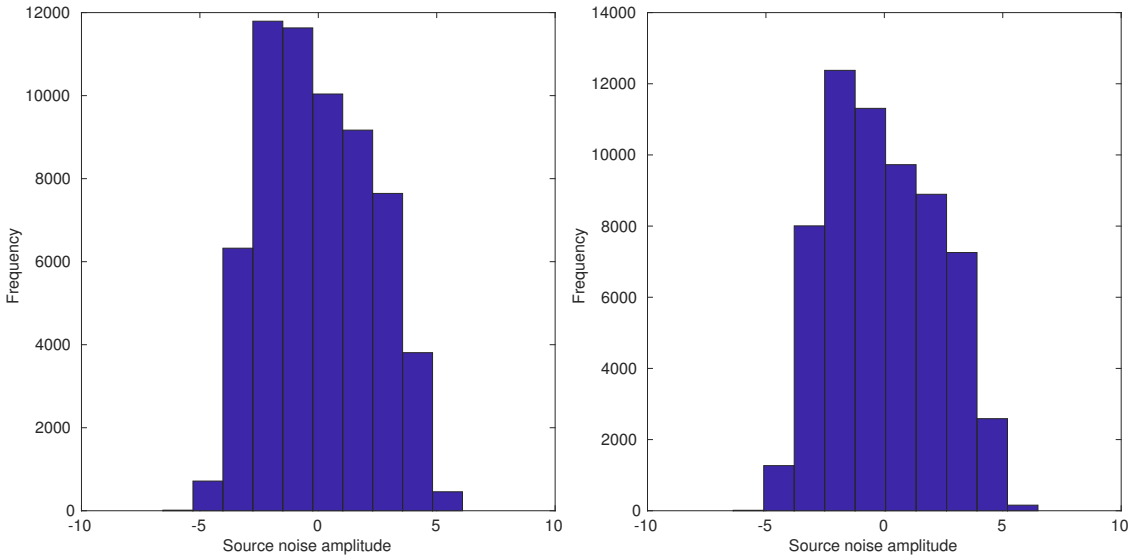

(a) Vertex 1599

(b) Vertex 10361

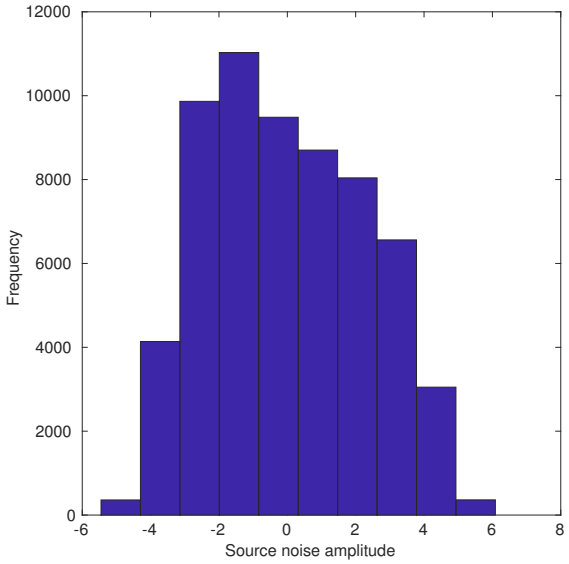

(c) Vertex 14965

Figure 3.13 Histograms of the source noise across time for three different vertices when the source noise was extracted using method 1.

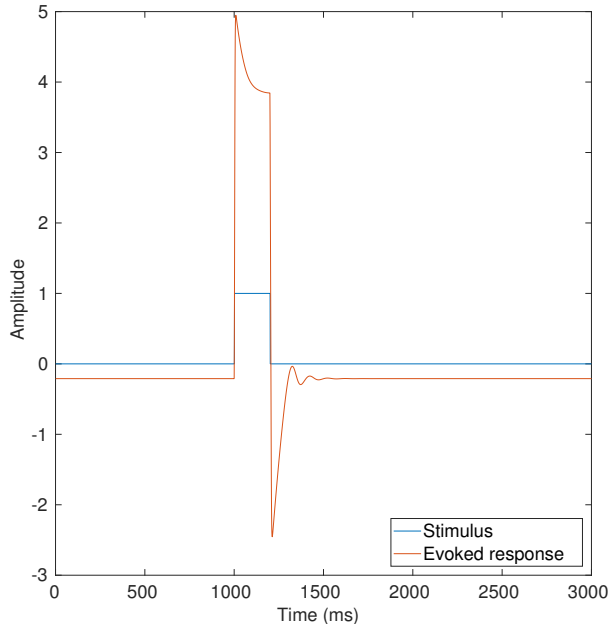

Figure 3.14 Comparison of a stimulus time-course and a stimulus evoked-response.

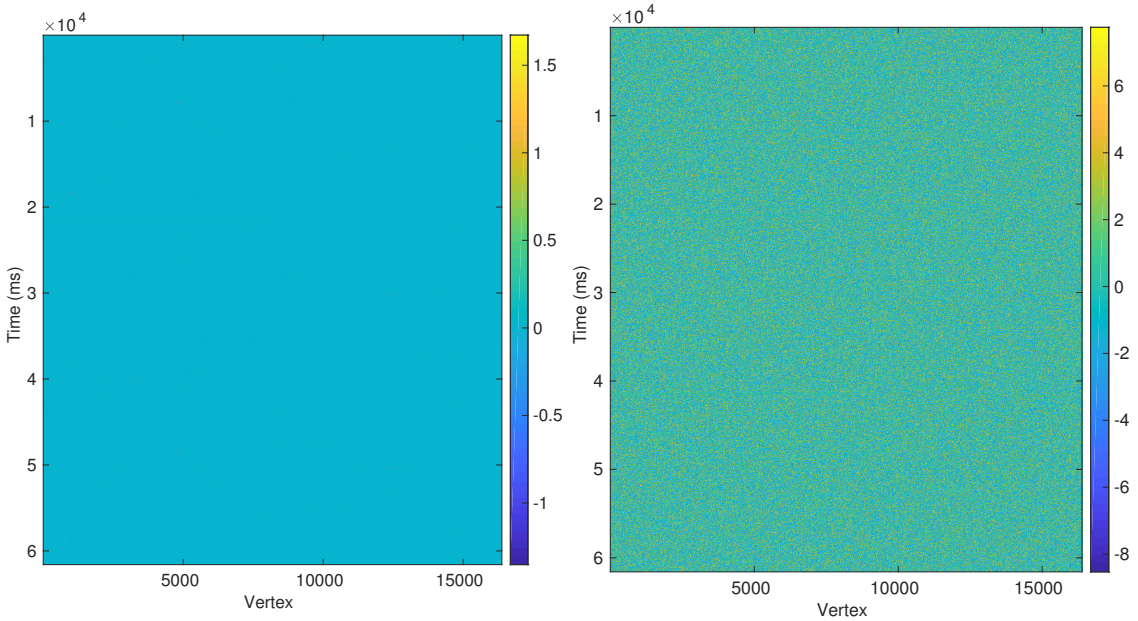

(a) Noise dispersion = 0 (no noise)

(b) Noise dispersion = 1

Figure 3.15 Source noise extracted using the stimulus-evoked response time-courses (method 2, see Section 3.3.2.1.2)

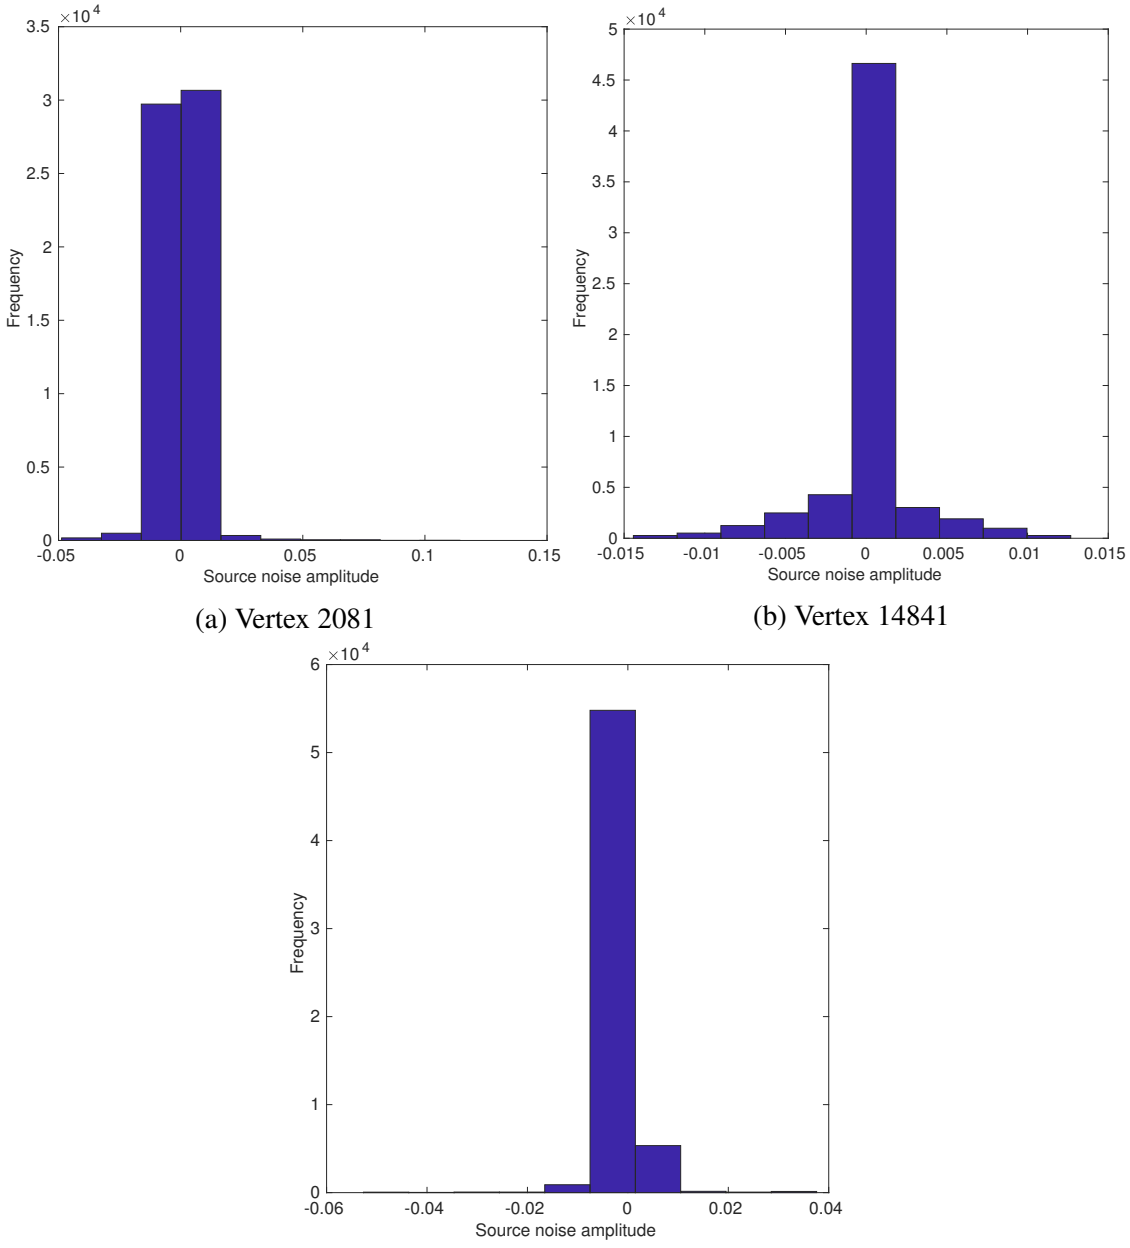

(a) Vertex 2081

(b) Vertex 14841

(c) Vertex 13349

Figure 3.16 Histograms of the source noise across time for three different vertices when source noise was extracted using method 2 and the noise dispersion = 0.

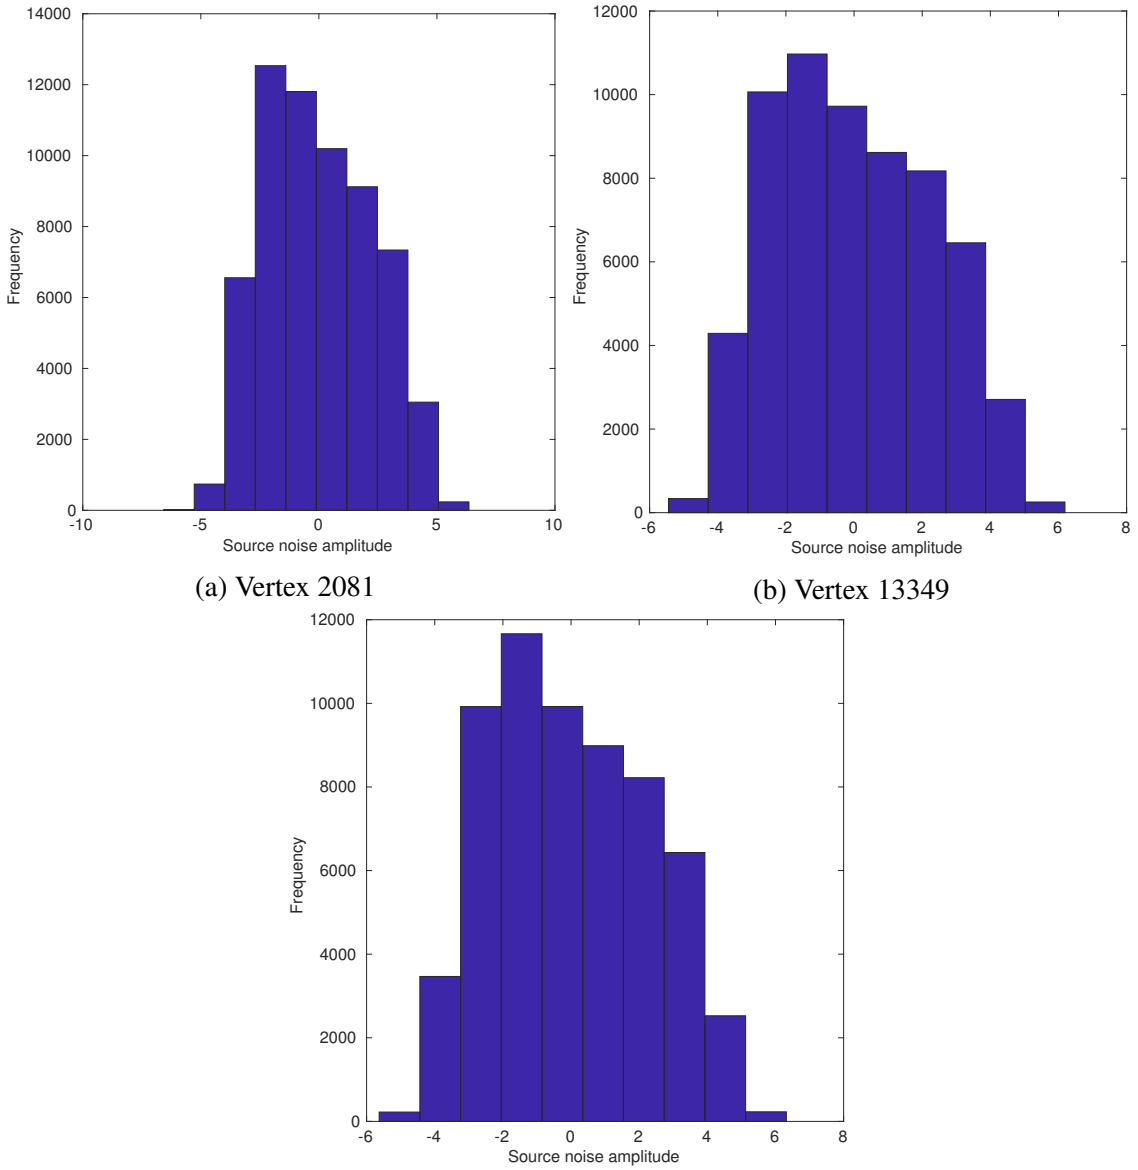

(a) Vertex 2081

(b) Vertex 13349

(c) Vertex 14841

Figure 3.17 Histograms of the source noise across time for three different vertices when source noise was extracted using method 2 and the noise dispersion = 1

3.3.2.2 Construction of predicted IC time-courses

To evaluate ICA, for each IC and source pair, a predicted IC time-course was constructed. To obtain this, the following model was first fit to each IC time-course using linear regression:

$$i(t) = \beta_0 + \beta_{s_1} s_1(t) + \beta_{s_2} s_2(t) + \cdots + \beta_{s_n} s_n(t) + \beta_{\varepsilon_1} \varepsilon_1(t) + \beta_{\varepsilon_2} \varepsilon_2(t) + \cdots + \beta_{\varepsilon_m} \varepsilon_m(t), \quad (3.6)$$

where $i(t)$ was the time-course of an IC, $s_k(t)$ was the stimulus time-course of the k-th stimulus, and $\varepsilon_j(t)$ was the time-course of the source noise at the j-th vertex. Then, for the IC and the j-th source, the predicted IC-time course was constructed as:

$$p_j(t) = \beta_0 + \beta_{s_j} s_j(t) + \beta_{\varepsilon_{j_1}} \varepsilon_{j_1}(t) + \beta_{\varepsilon_{j_2}} \varepsilon_{j_2}(t) + \cdots + \beta_{\varepsilon_{j_l}} \varepsilon_{j_l}(t), \quad (3.7)$$

where $\{j_1, \ldots, j_l\}$ were the indices of the l vertices that belonged to source j.

3.3.2.3 ICA evaluation

For each source, for each IC, the pairwise correlation between the predicted IC time-course and the IC time-course was calculated. For each source, the correlation with the highest absolute value was used as the performance measure. As mentioned in the introduction to this section, the basic idea behind this method was to evaluate each IC in terms of how much information it provided about a source in terms of the activations in response to its stimulus as well as the source noise of the vertices that belonged to it. This method was applied to the EEG signals simulated in Experiment 4 in Chapter 2 with the noise dispersion set to 1. ICA performance was only calculated for the first ICA decomposition (out of 30) in order to get preliminary results.

3.3.3 Results

A comparison of EEG ICA performance for the same ICA decomposition using the stimulus time-course method (the method used in Chapter 2) and the predicted IC time-course method is shown in Figure 3.18. Both methods of ICA evaluation showed the same trends, but the predicted IC time-course method always had higher performance.

For the first IC, the stimulus time-course of source 34 had the highest absolute β-value. For illustration, the IC time-course, and the stimulus and predicted IC time-courses for this source are shown in Figure 3.19. The correlation between the IC time-course and the stimulus time-course was $r_1 = 0.0455$ and the correlation between the IC time-course and the predicted IC time-course was $r_2 = 0.1221$.

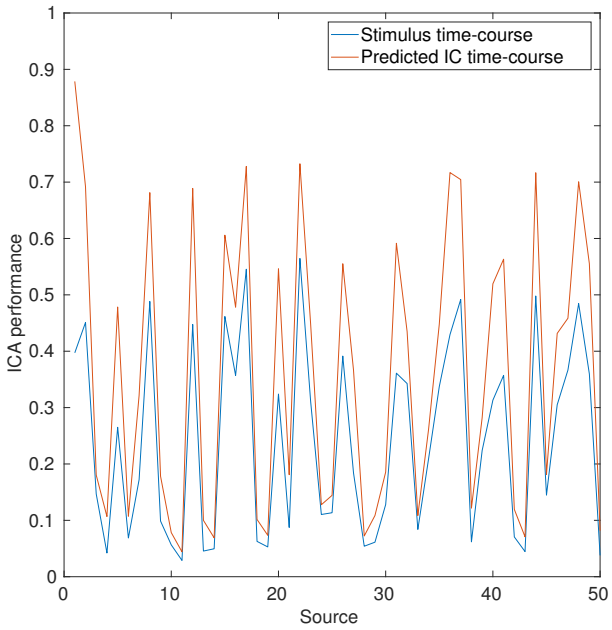

Figure 3.18 Comparison of ICA performance using the stimulus time-course method and the predicted IC time-course method.

3.3.4 Discussion

3.3.4.1 Interpretation

For an IC, the predicted IC time-course accounts for the source noise in the vertices belonging to the source for which its performance is evaluated against. The inclusion of the source noise improves the correlation with the IC suggesting that measuring ICA performance by including the source noise may be a better method than just using the stimulus time-course. The limitation of only using the stimulus time-course for evaluating ICA is that ICs are evaluated only in terms of how much information they contain about the stimulus-evoked activities (where the stimulus time-course is an approximation of the stimulus-evoked activities), and not how much information they contain about the source noise processes that are part of the source they correspond to. With the predicted IC time-course approach, for each IC, for each source, an estimation of how much the source noise of that source contributes towards the time-course of the IC is used in evaluating the correlation between the IC time-course and that of the source.

3.3.4.2 Over-fitting

The increase in performance when using the predicted IC time-courses instead of the stimulus time-courses might be a result of over-fitting as, while both methods evaluate the ICs in terms

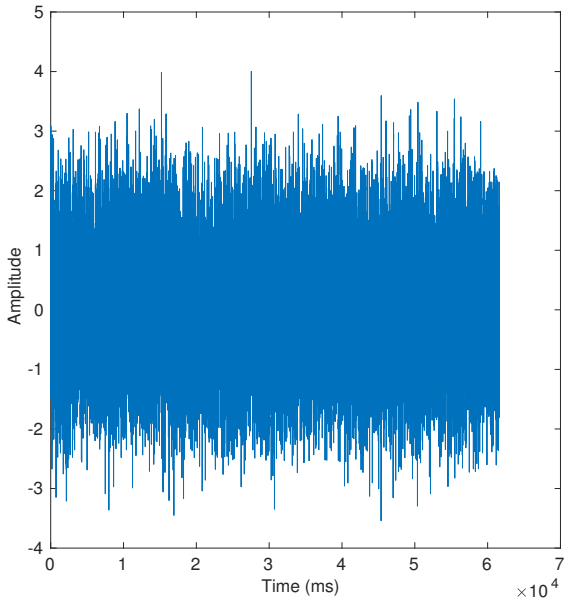

(a) IC time-course

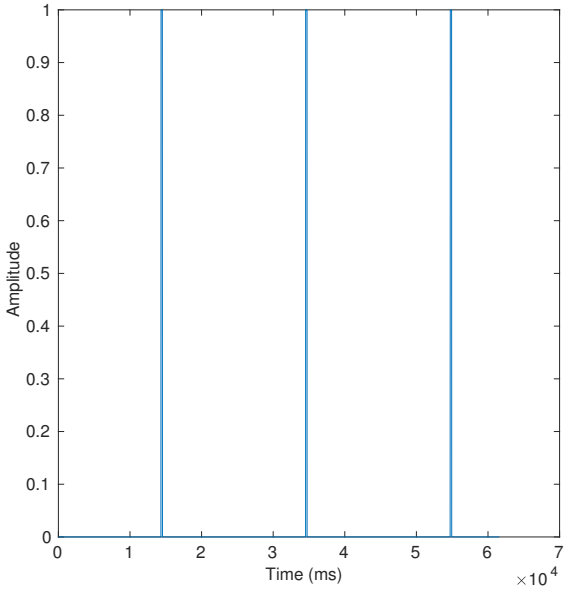

(b) Stimulus time-course

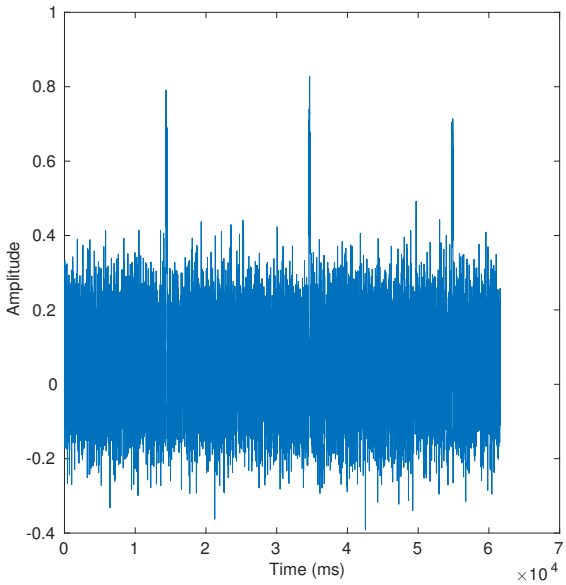

(c) Predicted IC time-course

Figure 3.19 Comparison of the stimulus time-course and predicted IC time-course for the first IC and the best matching source

of their correlation with the stimulus time-courses, the predicted IC time-course method also includes any variance within the IC time-course that could be explained by the source noise. Therefore, the correlation between the IC time-courses and the predicted IC time-courses will always be greater than or equal to that between the IC time-courses and the stimulus time-courses. One way to investigate if the inflation in performance is due to over-fitting or due to the IC time-courses actually containing noise information about the sources they correspond to could be to check the extent to which performance changes when the stimulus time-courses and the noise time-courses are mismatched in the predicted IC time-courses such that they correspond to different sources. Another way could be to inspect if ICs exist that have high correlations with the stimulus time-course of one source and the noise time-course of a different source. As it was decided that this method of evaluating ICA was infeasible (see Section 3.3.4.4), this was not investigated further as it was not going to be used in this work.

3.3.4.3 Extending to the evaluation of fMRI sICA

While the predicted IC time-course method was only applied on EEG signals and temporal ICA decompositions, it could be used for fMRI spatial ICA evaluation as well. This would be done by constructing a predicted IC time-course for each fMRI spatial IC using a method similar to that used for the EEG temporal ICs. First, the source noise time-courses would be convolved with the haemodynamic response function and then downsampled to match the sampling rate of the fMRI signals. Then, for each IC, the model described in Equation 3.6 would be fit with the regressors being the predicted fMRI signals of the sources and the convolved source noise time-courses. For each IC, for each source, a predicted IC time-course would be constructed by using Equation 3.7. The ICA evaluation would be performed by measuring the correlation between the spatial IC time-courses and the predicted IC time-courses.

3.3.4.4 Feasibility

For a single simulation, the source noise extraction took around 2 hrs 15 mins on a single node of a computing cluster with Intel Xeon E5 Family processors (12 to 24 cores) and either 96 GB or 128 GB of RAM. It also took approximately 100 minutes to fit the model described in Equation 3.6 to an IC. As there were 50 ICs in this simulation, it took 5000 minutes $\approx$ 84 hours in total to calculate ICA performance. Including the source noise extraction, the total time was $\approx$ 86 hours. For a single simulation, to run this process for all the 30 ICA

decompositions without parallel processing, it would have taken $30 \times 86 = 2580$ hours = 107.5 days.

3.3.5 Conclusions

ICA had better performance using the predicted IC time-courses for evaluation than just using the stimulus time-courses, indicating that the ICs did contain some relevant information about the source noise. This, however, could have been due to over-fitting. As the trends in performance between the two methods were similar and as this method was computationally infeasible, the accuracy of this method was not investigated further and it was not used for the evaluation of ICA.

3.4 Chapter summary

In this chapter, it was highlighted that the EEG and fMRI signals were not functions of only the stimulus time-courses and spatial maps respectively in the presence of source noise. In Chapter 2, both EEG temporal ICA and fMRI spatial ICA performances had been measured using the correlations between the IC time-courses and the stimulus time-courses. While the performance metrics in that chapter did indicate how much the information the ICs contained about stimulus-evoked activities (approximated by the stimulus time-courses), they did not indicate how much information was contained about the source noise, which, arguably, could also be considered as activity of interest as it pertains to source activity.

To investigate how to evaluate ICs in terms of how much source noise information they contained, along with information on stimulus-evoked responses, two different directions for developing the evaluation method were explored. In the first direction, the mean LFP signal across vertices in a source was used to approximate source time-courses, and in the second direction, predicted IC time-courses were used. The mean LFP approach for approximating the source time-courses was observed to not provide significant improvements over using the stimulus time-courses, and the reason for this was understood to be that this approach assumed that all the vertices had the same spatial location, which was incorrect. In the predicted IC time-course approach, the source noise of each vertex was first extracted and then, for each IC and source pair, a predicted IC time-course was constructed that contained variance explained by the source's stimulus time-course and the noise in the vertices it comprised. This method for evaluating ICA resulted in higher performance values than the method used in Chapter 2, but had the same trends. As this method was computationally infeasible, it was decided that

this method would not be used for further evaluation and, therefore, further investigation into whether its inflated performance was a result of over-fitting was not performed either.

One contribution of the work done in this chapter is that it highlights a potential limitation of the method used for evaluating ICA performance in Chapter 2 in that the evaluation was limited to only measuring how much information the ICs contained about the sources in terms of their stimulus-evoked activity and not their source noise. This, however, is only a limitation when the source noise is of interest, which was not the case in Chapter 2. Another contribution of the work done in this chapter is that it explores how ICA could be evaluated if the source noise were of interest and provides an initial investigation in this direction.

Chapter 4

Evaluation and Development of EEG-fMRI Parallel ICA

Abstract

Parallel independent component analysis (pICA) is a framework for combining information about latent sources from signals measured using multiple modalities. The basic idea is that independent component analysis (ICA) is first performed in each modality and then the independent components (ICs) are matched across modalities where they have statistical similarities. Together, the matched ICs are interpreted as providing multimodal source information. In the case of concurrently recorded electroencephalogram (EEG) and functional magnetic resonance imaging (fMRI) signals, the EEG IC provides a high resolution temporal description and the fMRI IC provides a high resolution spatial description. In this work, pICA is evaluated in terms of its accuracy and specificity at extracting information on underlying functional brain networks from concurrent EEG and fMRI signals. In contrast to how it has been used previously, it is applied here to the simulated data of a single subject without any temporal reduction to the signals or their IC time-courses, as this is the most basic form of the approach and should be used as the benchmark against which any further developments to it are compared. Two extensions to pICA that use spatial features and mutual information for matching the ICs are also demonstrated and evaluated. The results indicate that pICA performance increases with the number of orthogonal sources in the absence of noise, and decreases with the level of noise dispersion when the number of sources is fixed. The addition of spatial information improves pICA performance, whereas mutual information performs worse than correlation at matching the ICs. This work is the first evaluation of the EEG-fMRI pICA approach that uses time-courses for matching the ICs with simulated signals. It is also the first demonstration of how the The Virtual Brain (Sanz-Leon et al., 2013) can be used to validate the biological accuracy of an EEG-fMRI analysis method.

4.1 Introduction

Independent component analysis (ICA) is a technique for decomposing multivariate linear mixtures into individual components, where the latter are maximally statistically independent from each other. ICA has been applied extensively on electroencephalogram (EEG) and func-

tional magnetic resonance imaging (fMRI) signals for the purpose of extracting information on functional networks in the brain. On EEG signals, temporal ICA is typically applied, where the independent components (ICs) are separated from each other on the basis of their independence in the temporal domain, and on fMRI signals, spatial ICA, which optimises for spatial independence, is usually used. The main reason for this is that EEG signals have higher temporal resolution than spatial, and the opposite is true for fMRI signals. Since the early 1990s (Ives et al., 1993), concurrent EEG and fMRI recordings have been performed with the intention of combining the two modalities to obtain high resolution information on functional networks in both domains. To this end, ICA has been applied in various ways to extract this combined information. One of these methods is parallel ICA (pICA), which is the focus of this work. Here, a 'stripped-down' version of the pICA algorithm and some extensions to it are evaluated in terms of their accuracy using simulated EEG-fMRI signals.

ICA decomposes a matrix (Y), where each row of Y corresponds to a variable and each column to an observation, into the product of a mixing matrix (M) and a source matrix (S), i.e. $Y = MS$ (Hyvärinen and Oja, 2000). Each row of S is an IC and each column an observation corresponding to the same column in Y. Each column of M contains the weights with which the ICs are linearly mixed in Y. For temporal ICA, Y is arranged such that the rows correspond to spatial points (e.g. channels in EEG, voxels in fMRI) and the columns to time-points, and for spatial ICA it is arranged such that the rows are time-points and the columns are spatial-points. Consequently, for EEG temporal ICA, rows of S are IC time-courses and columns of M their corresponding scalp maps, and for fMRI spatial ICA, rows of S are IC spatial maps and the columns of M their corresponding time-courses. These are illustrated in Figure 4.1.

The basic idea of pICA for EEG-fMRI is to first perform ICA on the EEG data and the fMRI data separately and then to pair the independent components (ICs) across modalities where they have statistical similarities. EEG-fMRI ICs that are paired, also referred to as parallel ICs (pICs), are interpreted as providing information about the same source, i.e. functional network. There are two main approaches used when applying pICA. The first one can be termed as the 'matching across subjects' approach, as ICs that are similarly mixed across subjects in each modality are paired with each other. This approach was used when pICA was first proposed by Liu and Calhoun (2007). In this approach, for each subject, a summary endpoint for each modality is first obtained, such as an event-related potential (ERP) time-course from EEG and a statistical parametric map from fMRI. Then, for each modality, the data across subjects is vertically concatenated to construct Y on which ICA is subsequently performed. A modification is also made to the Infomax ICA algorithm (Bell and Sejnowski, 1995) where, for each modality, the optimisation is not just based on maximising

the independence between the components, but also on increasing the correlation between a single pair of components across the modalities. The output of this method is a single pair of EEG-fMRI ICs and the strength of the similarity between them is quantified by their correlation. In addition to EEG-fMRI, this approach has also been used to combine fMRI and single nucleotide polymorphism (SNP) data (Liu et al., 2009), and has been extended to three modalities (structural MRI, fMRI, and SNP) where a triplet of ICs, instead of a pair, are interpreted to provide information on the same source (Vergara et al., 2014).

The second pICA approach can be termed as the 'trial-by-trial' approach in which, rather than using a single summary endpoint per subject, the within-subject temporal dynamics are retained by applying ICA directly on the EEG and fMRI signals. Also, in contrast to the previous approach, the ICA decomposition in each modality is performed independently of the other. This was first used by Eichele et al. (2008) where EEG IC single-trial weights (trial-wise summary statistics) were convolved with a canonical haemodynamic response function (HRF) and then used as regressors in a linear model predicting the fMRI IC time-course for each fMRI IC. EEG and fMRI ICs were paired if their β-values were statistically significantly non-zero across subjects. This approach was also used by Eichele and Calhoun (2010) where, instead of convolving the EEG ICs with the HRF, deconvolution was used to extract the HRF from each fMRI IC and linear regression with the EEG IC single-trial weights as the dependent variable and the fMRI IC single-trial weights as the independent variable for each EEG-fMRI IC combination was used to pair the ICs where the β-values were significantly non-zero across subjects. Deconvolution has also been used on resting-state data where spatiospectral ICA (Y is structured such that each spatial-point has multiple columns for different frequency bands) is performed on the EEG signals, instead of temporal ICA, and the fMRI IC time-courses are deconvolved with respect to the EEG IC time-courses and the EEG and fMRI ICs are matched if the obtained impulse response functions are plausible HRFs (Bridwell et al., 2013; Wu et al., 2010).

The evaluation of EEG-fMRI analysis methods in terms of the biological accuracy of the information they provide is difficult to do with physically acquired signals as 'ground-truth' information, e.g. using intracranial EEG, cannot be obtained non-invasively. Therefore, the only readily available option is to use simulated data, for which the ground-truth can be known. The evaluation of pICA has only been performed using simulated data for methods that use the'matching across subjects' approach. For example, Liu and Calhoun (2007) evaluated both the *component accuracy* (accuracy of the ICs corresponding to sources within each modality) and the *connection accuracy* (accuracy of the connection strength between the matching EEG-fMRI IC pair) using simulated signals with 8 sources. Across multiple simulations, they varied the connection strength and the signal-to-noise ratio (SNR), and across ICA

decompositions, they varied the number of ICs. They also compared performance when the decompositions interacted across modalities compared to when they were independent. Overall, they observed that both methods had high component accuracy ($r > 0.9$), when the number of ICs was eight or more, and the 'interacting' method had slightly better connection accuracy, except when the ground-truth connection strength was less than 0.5, in which case it was inflated. In another example, Vergara et al. (2014) observed that the extension of pICA with three modalities performed more accurately than pICA between two modalities and single-modality ICA, and that its performance improved the stronger the connection strength between the modalities was.

The current work focuses on the evaluation and development of a pICA method that falls within the 'trial-by-trial' category, as this approach does not involve significantly reducing within-subject temporal information. Existing methods that use this approach, however, perform reduction at either the trial-level, by matching EEG-fMRI ICs using single-trial weights (Eichele and Calhoun, 2010; Eichele et al., 2008), or the epoch-level, by performing ICA on signals that have single values summarising activations within time windows (Bridwell et al., 2013; Wu et al., 2010). Another key element of all the pICA methods described is that they have only been applied to group data: in the 'matching across subjects' approach, EEG-fMRI ICs are combined based on how they are mixed across subjects, and in the 'trial-by-trial' approach, group ICA (Calhoun et al., 2001a, 2009) is first performed in each modality and the ICs are linked across modalities using their single-subject back-projections (see Eichele et al., 2008, for details). The simplest, 'stripped-down', pICA algorithm would be one that is applied on signals, without any temporal reduction, acquired from a single subject. As it provides the fundamental building block of 'trial-by-trial' pICA, this algorithm is evaluated in the current work. Apart from this, two extensions to pICA are also proposed and evaluated. In previous literature, EEG and fMRI ICs have only been matched using their temporal features, whereas here spatial features are used as well. Also, in previous literature, EEG and fMRI ICs are only matched using methods that look at linear associations between them, whereas here, the use of mutual information (MI) (Shannon, 1948), which can detect nonlinear associations, is demonstrated and compared with correlation.

Previously (Chapter 2), we evaluated ICA performance on EEG and fMRI signals separately in each modality as a function of the number of sources, the size of the sources, the level of source noise, and the duration of stimulation. In the current work, the performances of EEG ICA, fMRI ICA, and pICA, collectivity referred to as (p)ICA, are evaluated as a function of the number of sources (Experiment 1) and the level of source noise (Experiment 2). These parameters were selected as they were deemed to be most useful at providing a

first characterisation of pICA performance. One main difference between the previous work and this one is that the performance measure has been modified to account for the specificity of the ICs (further details in Section 4.2.5).

To summarise, this work is the first evaluation of pICA using the 'trial-by-trial' approach and it starts from the simplest algorithm, which is the building block of this approach, that is applied on the data of a single subject and does not perform temporal reduction on the EEG-fMRI signals or their IC time-courses. Apart from proposing a new method for evaluating (p)ICA, two extensions to pICA are demonstrated and evaluated. These are matching the EEG and fMRI ICs using (i) spatial features and (ii) MI.

4.2 Materials and methods

In the following subsections, The Virtual Brain (TVB) (Sanz-Leon et al., 2013) is first introduced as the simulator that is used to generate the EEG and fMRI signals on which pICA is evaluated. An overview of the pipeline for performing and evaluating pICA is described in Section 4.2.2, and the details of the steps are described in Sections 4.2.3 to 4.2.6. Experiment 1, which evaluates (p)ICA performance as a function of the number of sources is described in Section 4.2.7, and Experiment 2, which evaluates (p)ICA performance as a function of source noise dispersion is described in Section 4.2.8. In both of these experiments, pICA is evaluated using spatial features, which is a novel contribution of this work, along with temporal features, which have been used in previous literature (e.g. Eichele and Calhoun, 2010; Eichele et al., 2008). In Section 4.2.9, the method for extending pICA to use mutual information for matching the ICs across modalities and for evaluating its performance using this measure is described. This is also a novel contribution of this work.

4.2.1 The Virtual Brain (TVB)

The Virtual Brain (TVB) (Sanz-Leon et al., 2013) was used to simulate concurrent EEG and fMRI signals. In TVB, the human cortex is split into vertices, and sources are constructed by defining stimuli that evoke responses in the vertices. In this work, a vertex is referred to as a voxel, as it is used as the spatial unit of the fMRI signals. The stimuli are defined by their time-courses and spatial maps. All the simulations were constructed to have maximally orthogonal sources, so that pICA was evaluated in 'best-case' conditions. Each source comprised a single, unique region, that was randomly selected from a uniform distribution without replacement, and, therefore, there was no spatial overlap between different sources. The spatial map of a source had 1's for the voxels within the region it comprised, and 0's

for the others. The source time-courses were pulse-trains comprising three pulses of 0.2 s each, designed such that there was minimal overlap between their active periods by using $\theta_n = \theta_1 + (n-1)\Delta\theta$, where θ_n was the phase offset of the n-th source, $\Delta\theta = T/N$, $T = 20.2$ s was the interval between the pulses, N was the number of sources, and $\theta_1 = 3$ s.

Source noise is implemented in TVB using stochastic integration schemes (Sanz-Leon et al., 2015b). Heun's method for stochastic integration (see Appendix C of Sanz-Leon et al. (2015b)) was used to add white Gaussian noise to the source local field potential (LFP) signals that were then used to construct the EEG and fMRI signals via forward models. The noise processes had zero means and the noise dispersion was defined as their standard deviation. Within a simulation, the LFPs for all the voxels across the brain had the same level of noise added. The simulation parameters used and the source construction was the same as in our previous work (Chapter 2) and further details are provided there.

4.2.2 Pipeline

The steps involved in performing pICA and the evaluation are outlined in this subsection to provide an overview of the process. Details of each step are provided in the subsections that follow. The FastICA algorithm (Hyvärinen and Oja, 1997) was used to perform ICA using the implementation by Gävert et al. (2005).

Step 1 is to perform ICA in each modality. A single EEG temporal ICA decomposition and fMRI spatial ICA decomposition is shown for a simulation with 50 sources in Figure 4.1.

Step 2 is to measure the similarity between the EEG and fMRI ICs (see Section 4.2.3) and to keep the EEG-fMRI IC pairs between which the similarity exceeds the *pICA similarity threshold*. The unthresholded similarity between the previous EEG-fMRI IC pairs is shown in Figure 4.2.

Step 3 is to measure the similarity between the retained EEG-fMRI IC pairs and the sources (see Section 4.2.4) and to match the pairs and sources that have similarity that exceeds the *source similarity threshold*. An example of an EEG-fMRI IC pair and the source it corresponds to is shown in Figure 4.3.

Step 4 is to evaluate (p)ICA performance (see Section 4.2.5). An example of the performances across the pICA similarity threshold in the three pICA similarity domains is shown in Figure 4.4.

4.2.3 Measuring EEG IC and fMRI IC similarity

The similarity between the EEG and fMRI ICs was measured in three *pICA similarity domains*: temporal, spatial, and spatiotemporal. In the temporal domain, the EEG IC time-

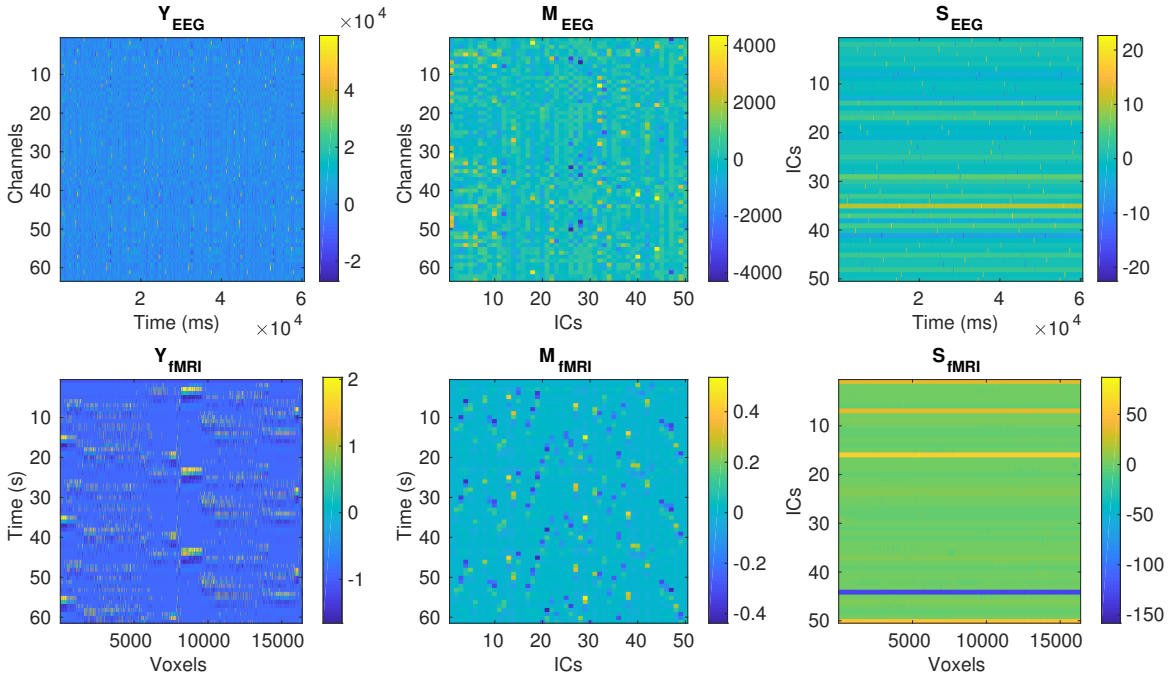

Figure 4.1 A single EEG temporal ICA decomposition (top row) and fMRI spatial ICA decomposition (bottom row) for a simulation. Top row: Y_{EEG} is the EEG data, where rows are channels and columns time-points (sampled at 1000 Hz). S_{EEG} rows are IC time-courses, and the columns in M_{EEG} are their corresponding scalp maps. Bottom row: Y_{fMRI} is the fMRI data, where rows are time-points (sampled at 1 Hz) and columns are voxels. S_{fMRI} rows are IC spatial maps, and the columns in M_{fMRI} are their corresponding time-courses.

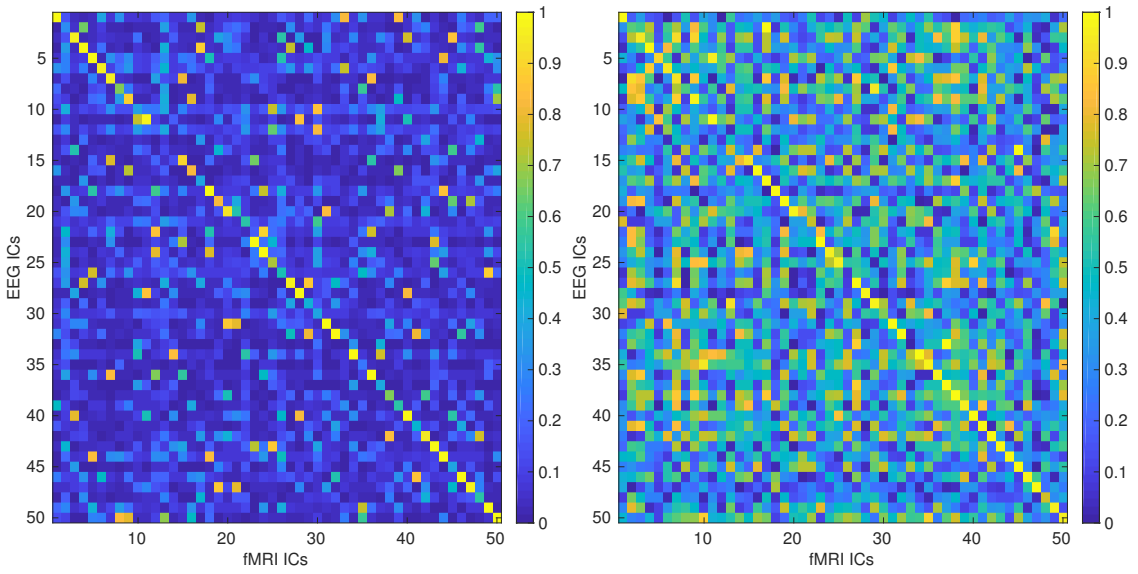

(a) Temporal (b) Spatial

Figure 4.2 Similarity between the EEG-fMRI ICs shown in Figure 4.1 in the (a) temporal and (b) spatial domains. (Note: the fMRI ICs have been sorted such that the maximum value of each EEG IC lies on the diagonal. In cases where multiple EEG ICs had maximum values with the same fMRI IC, the fMRI IC was positioned such that the maximum value of the bottom-most EEG IC was on the diagonal.)

courses were convolved with the first order Volterra series HRF (Friston et al., 2000) that was used to simulate the fMRI signals and downsampled to 1 Hz so that they matched the temporal resolution of the fMRI ICs, and then the magnitude of the Pearson's correlation between the IC time-courses was used as the similarity measure. In the spatial domain, the fMRI IC spatial maps were projected to the scalp using the projection matrix provided in TVB, and then the magnitude of the Pearson's correlation between the IC scalp maps was used as the similarity measure. In the spatiotemporal domain, the similarity measure was the mean of the temporal and spatial similarity measures. In the investigation with mutual information (MI) (Section 4.2.9), MI was used instead of the magnitude of the Pearson's correlation.

4.2.4 Measuring EEG-fMRI IC pair and source similarity

The similarity between the EEG-fMRI IC pairs and the sources was measured in four *source similarity domains*: temporal, spatial, spatiotemporal, and modality-specific. In all cases, the similarity was calculated as the mean of the similarity across the two modalities. In the temporal domain, similarity for the EEG IC was calculated as the magnitude of the

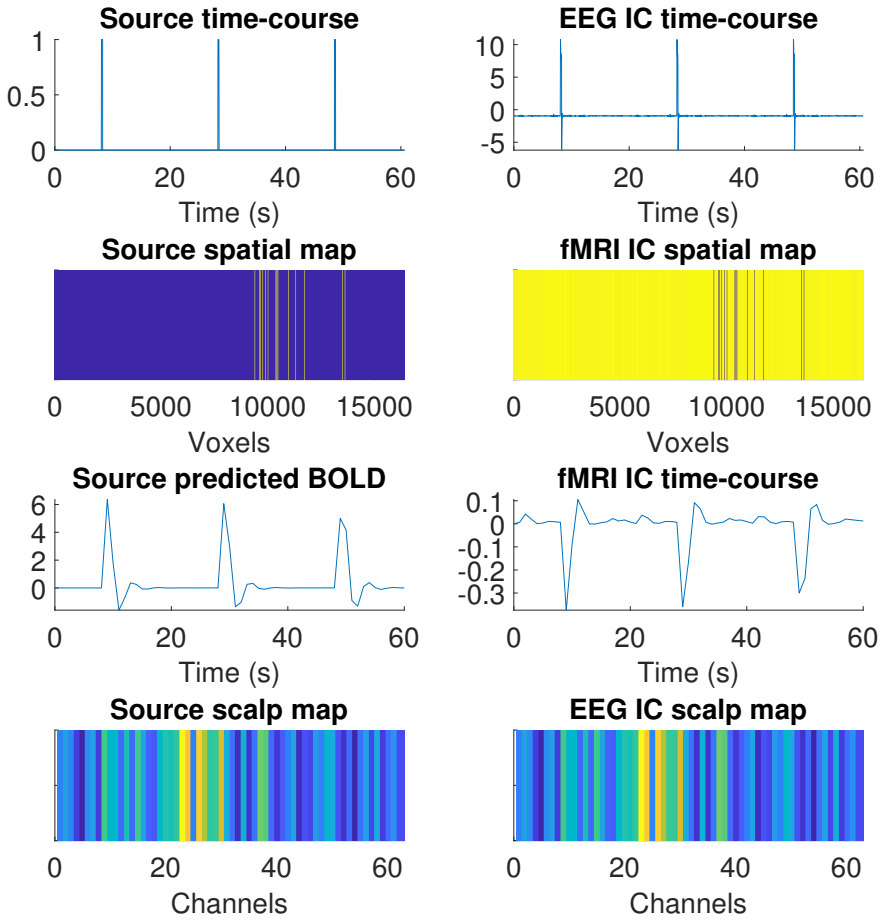

Figure 4.3 Example of a source (left column) and its corresponding EEG-fMRI IC pair (right column) with a similarity of 0.981 in the temporal domain. The EEG-fMRI IC pair provides an accurate description of the source in the temporal (EEG IC time-course) and spatial (fMRI IC spatial map) domains.

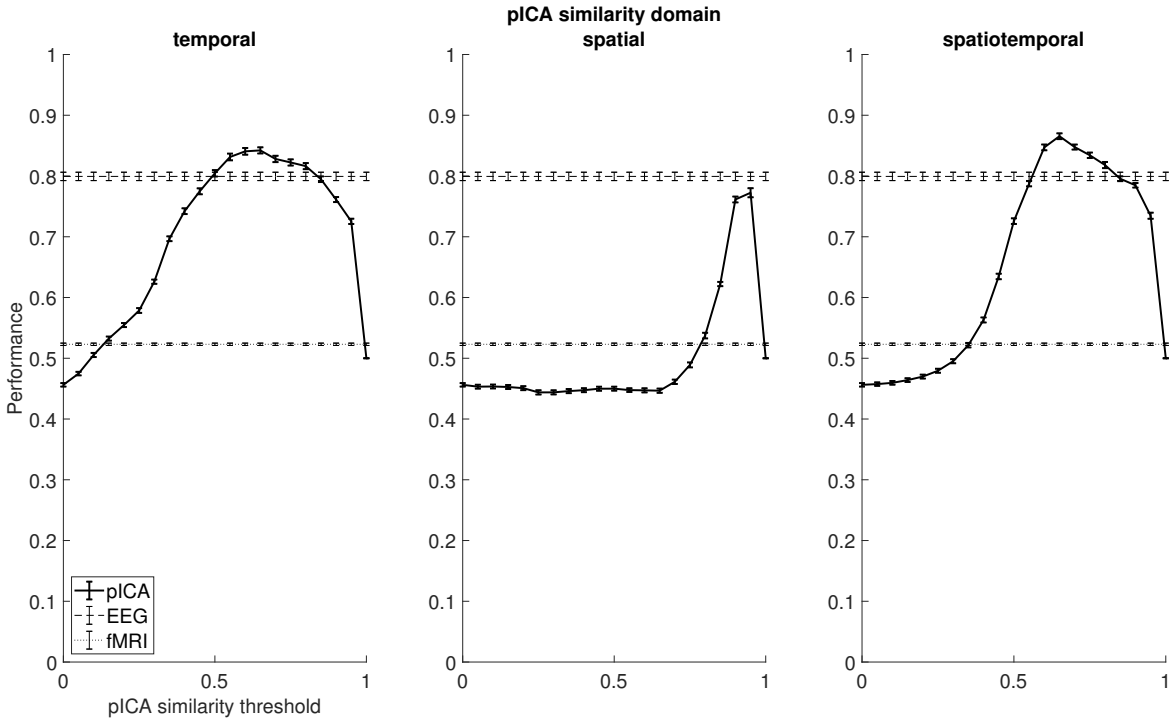

Figure 4.4 Example of (p)ICA performance across the pICA similarity threshold in the three pICA similarity domains (temporal, spatial, spatiotemporal). The EEG and fMRI ICA performances are constant as they do not depend on the pICA similarity domain or threshold. The source similarity domain was temporal and the source similarity threshold was 0.8. The performance measures were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions and the errors bars to the standard error of the mean.

Pearson's correlation between the IC time-course and the source time-course, and similarity for the fMRI IC was calculated as the magnitude of the Pearson's correlation between the IC time-course and the predicted BOLD time-course of the source, which was constructed by convolving the source time-course with the first order Volterra series HRF (Friston et al., 2000) that was used to simulate the fMRI signals and downsampling it to 1 Hz. In the spatial domain, similarity for the fMRI IC was calculated as the Pearson's correlation between the IC spatial map and the source spatial map, and similarity for the EEG IC was calculated as the magnitude of the Pearson's correlation between the IC scalp map and the source scalp map, which was constructed by multiplying the EEG projection matrix with the source spatial map. In the spatiotemporal domain, similarity for the EEG and fMRI ICs was calculated as their mean similarity across the temporal and spatial domains. In the modality-specific domain, similarity for the EEG IC was calculated in the temporal domain and similarity for the fMRI IC was calculated in the spatial domain, i.e. the domain in which the modality had higher resolution was used. In the investigation with MI (Section 4.2.9), MI was used instead of the magnitude of the Pearson's correlation.

4.2.5 Performance measures

For a simulation, three performance measures were calculated. These corresponded to pICA, EEG ICA, and fMRI ICA. Each of these measures was calculated using the same method. First, the following two quantities were calculated:

1. Proportion of sources with at least one matching (p)IC out of the total number of sources.

2. Proportion of (p)ICs with exactly one matching source out of the total number of (p)ICs.

The first quantity indicates how well (p)ICA performed at providing source information across the range of sources. The second quantity indicates the specificity of (p)ICs in terms of providing information about single sources. As the relationship between (p)ICA performance and both of these quantities is in the positive direction, and as the quantities have the same range (0 to 1), (p)ICA performance was calculated as their mean.

In cases where there were no pICs, which would happen when the similarity between any of the EEG and fMRI ICs did not exceed the pICA similarity threshold, the value of the first quantity was set to 0 and the second quantity to 1. The rationale behind this was to keep these two measures monotonic with respect to the number of pICs: the first quantity decreases as the number of pICs decreases, and so it was set to 0 when there were no pICs,

whereas the second quantity increases as the number of pICs decreases, and so it was set to 1 when there were no pICs. As a consequence, in these cases, pICA performance was 0.5.

Previously, we evaluated single-modality ICA performance using a slightly different performance measure (Chapter 2). The differences between the previous measure and current measure are: (i) the previous measure accounted for the performance of a single-source, whereas the current one accounts for all the sources within a simulation, (ii) the previous measure was the correlation between the source and the IC it had the highest correlation with, whereas the current measure is based on a classifier that decides whether or not a source matches an IC based on the source similarity threshold, (iii) the previous measure was only calculated in the temporal source similarity domain, whereas the current one is calculated in the spatial and spatiotemporal domains as well, and (iv) the previous measure did not account for the specificity of the ICs, which the current one does using the second quantity. The previous measure was appropriate to use in Chapter 2 as some of the experiments looked at how ICA performance varied as a function of the size of a source (the number of regions it comprised) and for this manipulation only the size of a single source was varied across simulations and ICA performance for this particular source measured. The current measure, however, is a progression from our previous work, with the main improvement being accounting for IC specificity.

4.2.6 Choosing the source similarity domain and threshold

To choose the source similarity domain and threshold to use for evaluating (p)ICA in Experiment 1 (Section 4.2.7) and Experiment 2 (Section 4.2.8), the performance of (p)ICA was evaluated on a single simulation across the domains and the range of the threshold. In this simulation, there were 50 sources with orthogonal time-courses and spatial maps and no source or sensor noise.

To first inspect if there was much variance in the similarity between the EEG-fMRI IC pairs and sources across the domains, unthresholded source similarity matrices for each EEG-fMRI IC pair were visually inspected. The domains were judged to be consistent in terms of the sources and EEG-fMRI IC pairs that had high similarity. The matrices are shown in the supplementary material in Figure C.1.

To further compare the domains and to identify the optimum source similarity threshold, (p)ICA performance across the source similarity threshold was plotted for each domain when the pICA similarity threshold was 0 (Figure 4.5a) and when it was 0.5 (Figure 4.5b) in the temporal pICA similarity domain. The former case was investigated as it looked at pICA performance considering all possible EEG-fMRI IC pairs and hence corresponded to the matrices previously inspected, and the latter case was investigated as, in practice, pICA

similarity would always be thresholded. This was done in the temporal domain, as that is what has been previously used in literature, and the threshold of 0.5 was deemed to be an appropriate arbitrary threshold for this purpose.

EEG performance was high in the temporal domain and low in the spatial, and the opposite effect was seen for fMRI. This was expected as EEG has low specificity in the spatial domain (many ICs match the same source) and fMRI has low specificity in the temporal domain. The performance of pICA had a peak in all cases, and the curves were similar for the two different pICA similarity thresholds in all domains except temporal. In this domain, pICA performance had significantly larger area under the curve, indicating that pICA performed significantly better when there were fewer EEG-fMRI IC pairs. This was expected, to some degree, in each domain as the second quantity described in Section 4.2.5 is higher with fewer pairs. It is unclear why there was a greater change in the temporal domain compared to the others. This was not further investigated as it was not considered to be important in deciding which source similarity domain to use.

Based on these initial investigations, it was speculated that the choice of the source similarity domain would not result in significant variation in (p)ICA performance, and therefore rigourous analysis to inform this decision was not needed. It was decided to use the spatiotemporal domain for the (i) a priori reason that it is unbiased in the selection of spatial and temporal features for both modalities, and the (ii) observed reasons that both EEG and fMRI performed well in this domain for source similarity thresholds greater than 0.5 and the peak was broader than in the modality-specific domain, indicating there was more margin for error in the selection of the source similarity threshold.

As it was not expected that different simulations would have the same performance profile as the one used to compare the source similarity domains, and as pICA performance had a narrow peak in all cases, it was decided to not use a single source similarity threshold, but rather a range that varied from 0.5 to 0.8 in steps of 0.1. In the main text of this article, the results with the threshold set to 0.8 are shown, and the others are shown in the supplementary material. It was decided to mainly report results with the threshold set to 0.8 for the a priori reason that 0.8 is considered to indicate strong correlation and, therefore, the source to (p)IC mapping with this threshold would be mostly (and probably entirely given the orthogonal design of the sources) accurate.

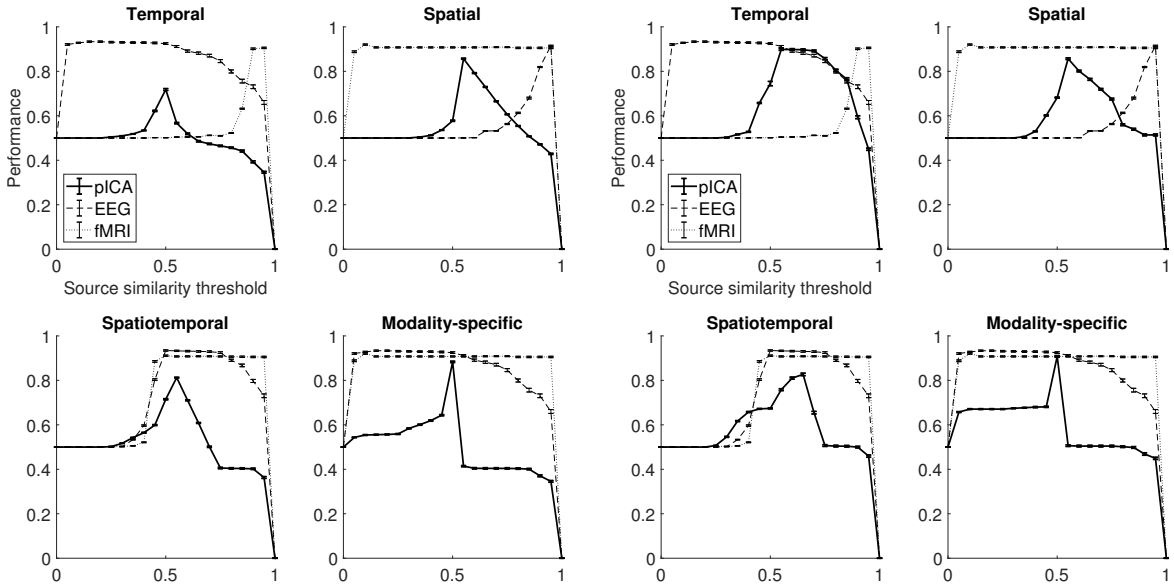

(a) pICA similarity threshold = 0 (b) pICA similarity threshold = 0.5

Figure 4.5 (p)ICA performance for a simulation with 50 orthogonal sources in each of the source similarity domains across the source similarity threshold with the pICA similarity threshold set to (a) 0 and (b) 0.5. The performance measures were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions and the errors bars to the standard error of the mean.

4.2.7 Experiment 1: Evaluation of (p)ICA performance as a function of the number of sources

To investigate how (p)ICA performance varied as a function of the number of sources, the EEG-fMRI signals previously simulated using TVB (see Experiment 1 in Chapter 2) were used. Sixty-one simulations were run and, across them, the number of sources was varied from 1 to 61. The maximum number of sources was set to 61 as the simulated fMRI signals had 61 time points. As described in Section 4.2.1, the sources were maximally orthogonal so that (p)ICA was evaluated in 'best-case' conditions. There was no source or sensor noise.

The purpose of this experiment was to explore how (p)ICA performance varied as a function of the number of sources across the pICA similarity domains and thresholds. The motivation behind this was (i) to check if in any domain, for any threshold, for any number of sources, (p)ICA provided biologically accurate source information and (ii) if it did, to identify what these parameters were. Based on previous work evaluating ICA within each modality (Chapter 2), it was expected that fMRI ICA performance would be high and largely unaffected by the number of sources whereas EEG ICA performance would linearly increase. However, in Chapter 2, the performance measure had not accounted for the specificity of the ICs, and therefore it was also expected that the previous results may not be replicated. The performance of pICA was expected to be high when both EEG and fMRI ICA performed well. As EEG ICA performance was expected to linearly improve with the number of sources, pICA performance was also expected to follow the same trend.

4.2.8 Experiment 2: Evaluation of (p)ICA performance as a function of source noise dispersion

To investigate how (p)ICA performance varied as a function of the level of additive, white Gaussian source noise, 21 simulations were run across which the noise dispersion (standard deviation of the noise process) was varied from 0 to 0.2 in steps of 0.01. All of these simulations had the same maximally orthogonal 50 sources. The hypothesis was that (p)ICA performances would decrease with noise dispersion. The motivation behind this experiment was to identify the extent to which pICA is robust to source noise by checking for what levels of dispersion it performed well (> 0.8).

To illustrate the effect of source noise, Figure 4.6a shows the time-course of a voxel with different levels of noise dispersion. The signal-to-noise ratio (SNR) for the voxel is shown in Figure 4.6b. The 'signal' power was calculated as the square of the root mean square (RMS) of the clean signal (dispersion = 0) and the noise power for each dispersion was calculated as the square of the RMS of the difference between the noisy signal and the clean signal.

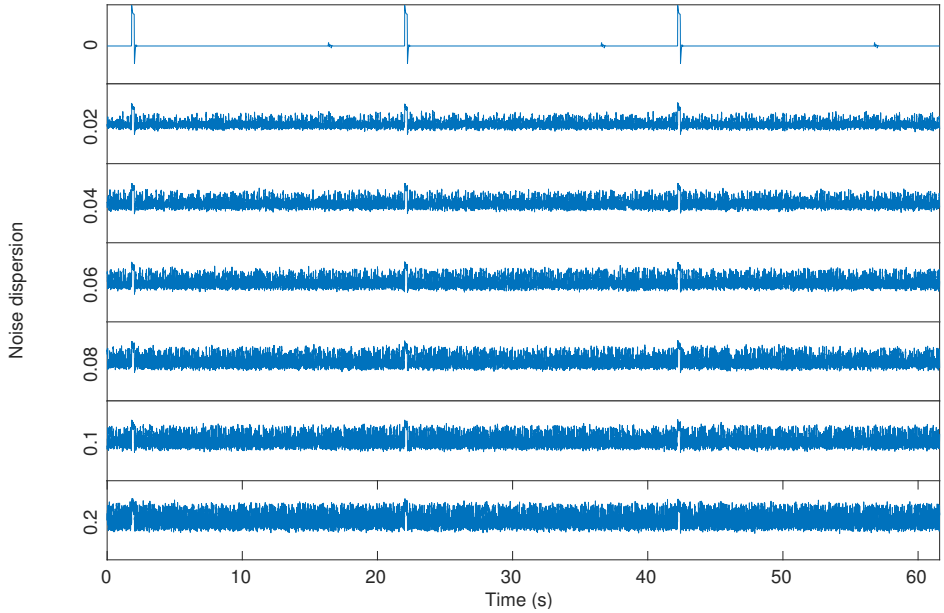

(a)

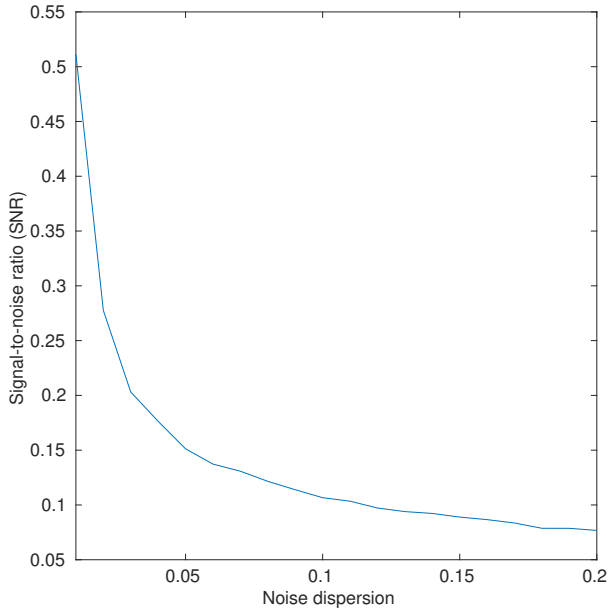

(b)

Figure 4.6 The source noise dispersion was varied from 0 to 0.2 in steps of 0.01. For a single voxel (a) shows the local field potential (LFP) time-courses and (b) the signal-to-noise ratio (SNR) across the different noise levels (excluding noise dispersion = 0 as in this case SNR $\rightarrow \infty$).

In previous work (Chapter 2) we evaluated the performance of ICA in each modality across noise dispersion of 0 to 2. In the current work, initially the experiment was performed with this range of dispersion (see Section C.4 in the supplementary material). Based on the results, it was concluded that the SNR was too low to draw any inference on pICA performance in the presence of noise. This, however, was not considered to question the validity of the inferences drawn in Chapter 2 as, in that work, it was expected that the 'sensor-level' SNR (opposed to the 'voxel-level' SNR shown in Figure 4.6a) would have been higher, as ICA performance was evaluated for a single source that had a size of 37 regions instead of a single region, which was the case in the current work.

4.2.9 Testing mutual information (MI) for measuring pICA similarity

EEG and fMRI ICs have only been matched in pICA using measures of linearity. To investigate how mutual information (MI) (Shannon, 1948), which can measure nonlinear relationships, performed as a pICA similarity measure, pICA was applied and evaluated on a single simulation comprising 50 orthogonal sources and no source or sensor noise. This investigation was exploratory to see if there were differences in pICA performance between the two measures and, if so, to inform whether pICA should be further developed to incorporate nonlinear measures. As the standard errors in pICA performance across multiple ICA decompositions in Experiments 1 and 2 had been small, this investigation was only done with a single EEG ICA and fMRI ICA decomposition. MI was calculated using MutualInfo 0.9 (Peng et al., 2005). It was performed on the z-scores of the variables as that significantly reduced the computation time.

The pICA similarity matrices in the temporal and spatial domains for both measures for a single EEG ICA and fMRI ICA decomposition are shown in Figure 4.7, and a scatterplot of their relationship, in the spatiotemporal domain, is shown in Figure 4.8. The Pearson's correlation between the two is 0.41. In Figure 4.8, it is clear that there are some cases where one measure shows EEG-fMRI IC pairs to be strongly related and the other does not (top-left and bottom-right regions). Figure 4.9 shows the three cases circled in Figure 4.8. In case (a), both MI and correlation have high similarity (MI = 1.22 bits, $r = 0.99$), and it is clear that the EEG and fMRI ICs have almost identical scalp maps and time-courses (the EEG IC predicted BOLD time-course is compared with the fMRI IC time-course, see Section 4.2.3). In cases (b) and (c), there are mismatches between the two measures, with (b) having high MI and low correlation (MI = 1.04 bits, $r = 0.04$) and (c) having low MI and high correlation (MI = 0.12 bits, $r = 0.98$). In case (c), based on visual inspection, it is clear that the EEG and fMRI IC scalp maps and time-courses are almost identical, and, therefore, in this case, correlation is more accurate than MI in matching the EEG-fMRI IC pair. Case (b), however, is less

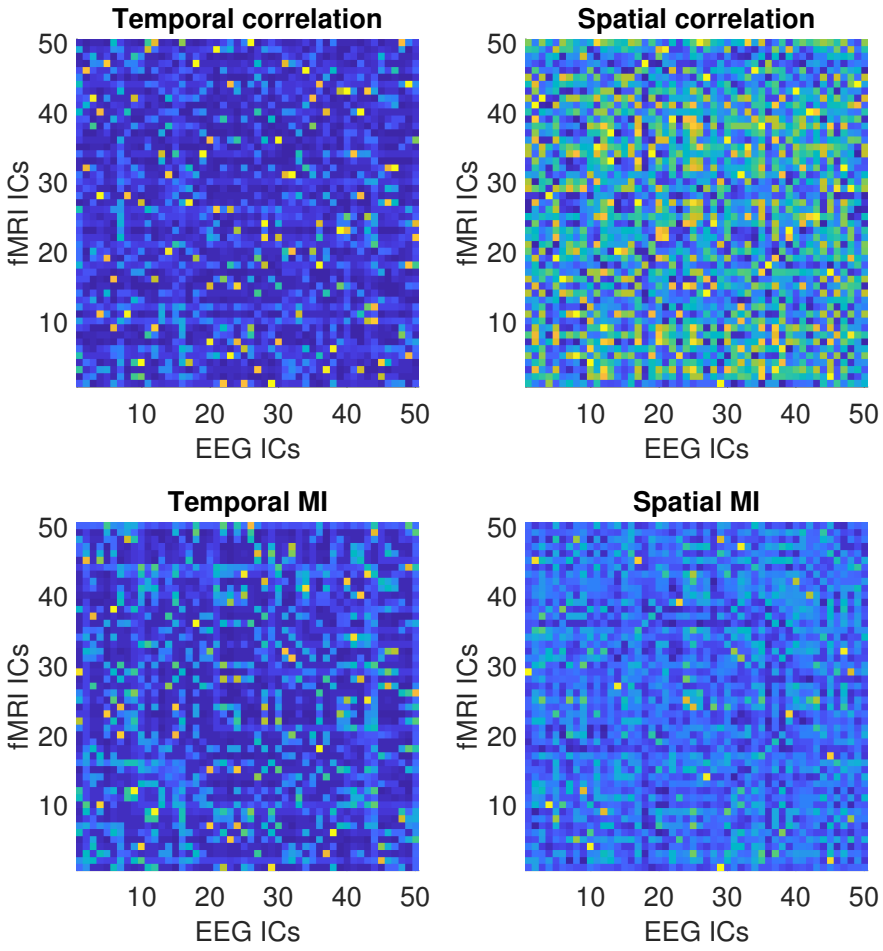

Figure 4.7 Similarity between the EEG-fMRI IC pairs when the pICA similarity measure was correlation (top row) and MI (bottom row) and the pICA similarity domain was temporal (left column) and spatial (right column). A single EEG ICA and fMRI ICA decomposition was performed.

clear. Further inspection of the IC time-courses (Figure C.16 in the supplementary material) shows that the peaks in the EEG predicted BOLD time-course are earlier than that in the fMRI IC time-course, and, therefore, in this case as well, correlation is more accurate than MI as it does not match the pair. Based on these initial observations, it was concluded that, in the cases in which there were mismatches between correlation and MI, correlation was more accurate at pairing the ICs. Other cases were not inspected, and further investigation qualitatively comparing the accuracy of the two measures was not performed, as this was considered beyond the scope of the current work, which focuses on performing an initial inspection on the performance of pICA with MI as the pICA similarity measure.

The comparison of MI with correlation as a pICA similarity measure poses the question on which measure should be used as the source similarity measure. This cannot be determined

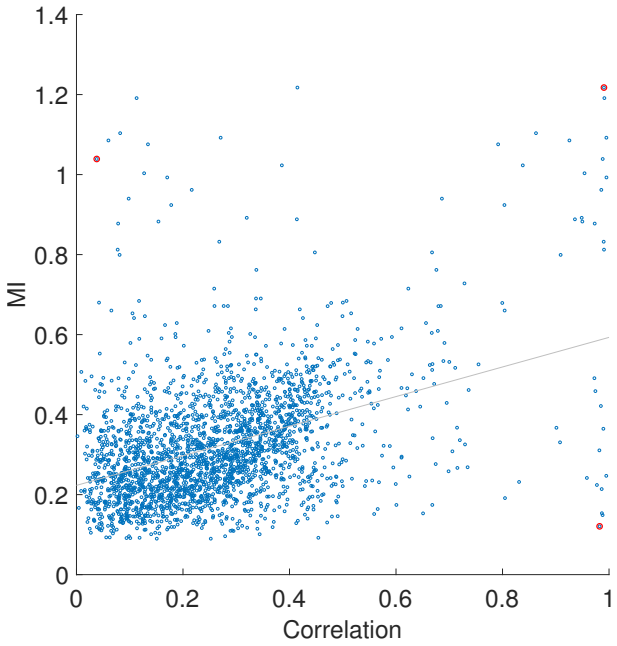

Figure 4.8 Comparison of the pICA similarity between EEG-fMRI IC pairs when the pICA similarity measure was correlation and MI in the spatiotemporal pICA similarity domain ($r = 0.41$). A single EEG ICA and fMRI ICA decomposition was performed. The three pairs identified by the red circles are shown in Figure 4.9

without ground-truth knowledge on which measure more accurately matches ICs to sources. As this knowledge is unavailable, it would be unfair to evaluate MI as a pICA similarity measure using correlation as the source similarity measure, and vice versa, as either measure may pick up valid associations between sources and ICs that the other is blind to. To first inspect how different the source similarity matrices were when using correlation or MI, a scatterplot of the unthresholded similarities between all sources and all EEG-fMRI IC pairs of both measures in the spatiotemporal domain was generated (Figure 4.10). While, in this domain, both measures were strongly correlated ($r = 0.84$), the relationship between them was not simply linear in the temporal and spatial domains (Figure C.17 in the supplementary material); for example, in the spatial domain, there were two clusters separated when correlation was around 0.5.

From these figures it was inferred that the two measures did not have similar source similarity matrices. Therefore, pICA performance was evaluated with MI as the pICA similarity measure and with both, MI and correlation, as the source similarity measures. For comparison, the same analysis was performed with correlation as the pICA similarity measure as well. The pICA and source similarity thresholds for both measures were from 0 to 1 in steps of 0.1. This was determined to be a suitable range for MI from the values in

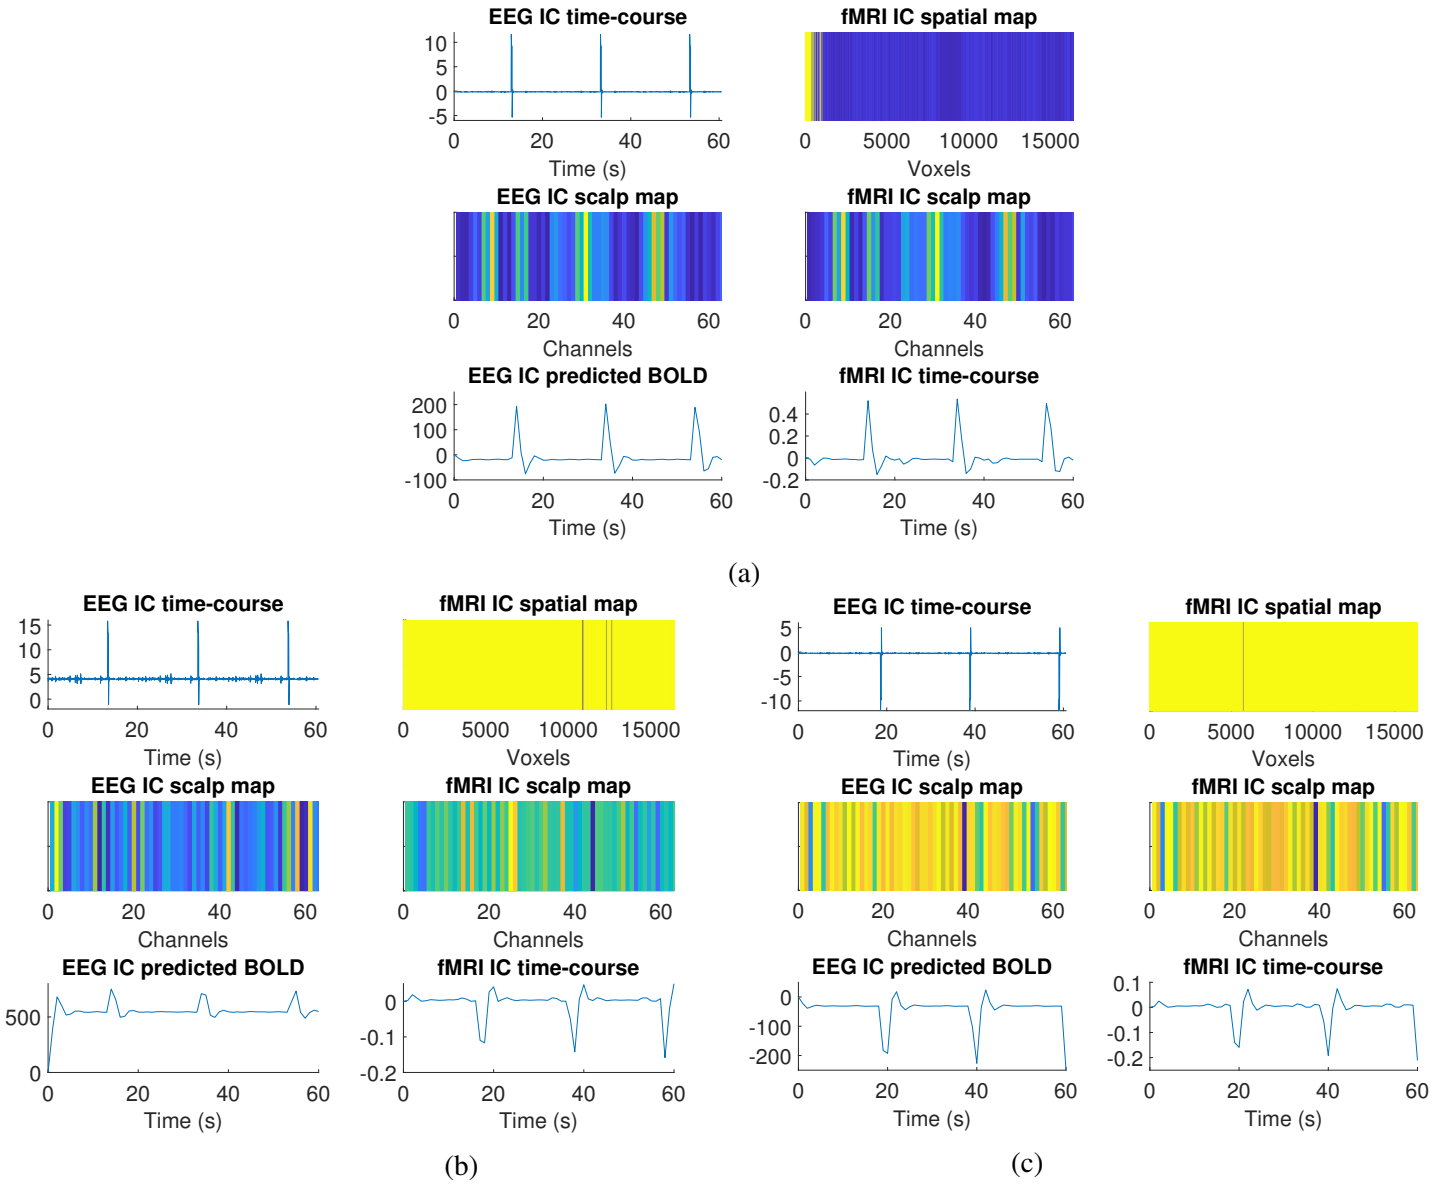

Figure 4.9 The three pairs of EEG-fMRI ICs marked on Figure 4.8. (a) high MI (1.22 bits), high correlation ($r = 0.99$), (b) high MI (1.04 bits), low correlation ($r = 0.04$), (c) low MI (0.12 bits), high correlation ($r = 0.98$). For each pair, the EEG IC features (left column) and fMRI IC features (right column) are shown. As the pICA similarity domain is spatiotemporal, the similarity between the ICs is the mean of the similarity between their EEG and fMRI IC scalp maps and their EEG IC predicted BOLD and fMRI IC time-courses.

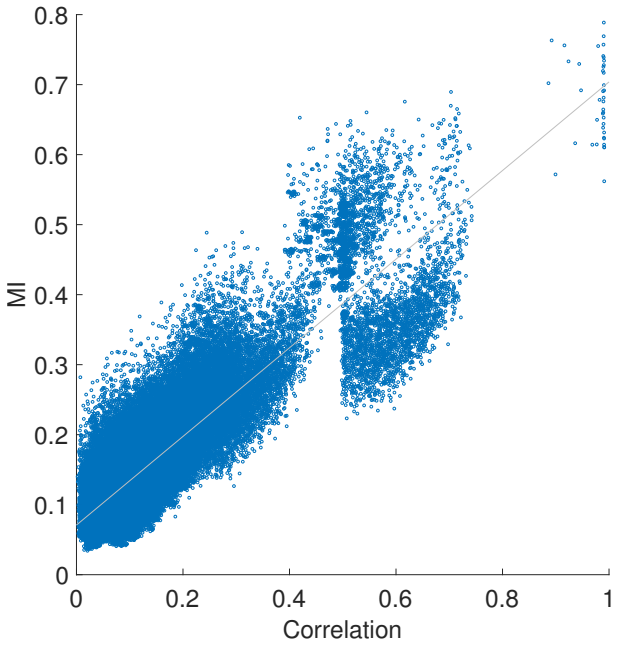

Figure 4.10 Comparison of the source similarity matrices when using correlation vs. MI as the source similarity measure in the spatiotemporal source similarity domain. The scatterplot shows the similarity between each source and EEG-fMRI IC pair ($r = 0.84$). These were obtained from a single EEG ICA and fMRI ICA decomposition on signals from a simulation with 50 orthogonal sources with no source or sensor noise.

Figure 4.10. Based on the the results of the previous experiments, both the pICA and source similarity domains were spatiotemporal. To identify the extent to which pICA performance was determined by ICA performance within each modality, EEG and fMRI ICA performances for each source similarity measure across the source similarity threshold were also evaluated.

4.3 Results

4.3.1 Experiment 1: (p)ICA performance as a function of the number of sources

The performance of (p)ICA as a function of the number of sources is shown in Figure 4.11. The source similarity domain was spatiotemporal and the source similarity threshold was 0.8 (see Section 4.2.6). Each column corresponds to a pICA similarity domain (temporal, spatial, spatiotemporal) and each row to a pICA similarity threshold (0.5 to 0.8, step 0.1). For each simulation, ICA was performed 30 times and the curves correspond to the mean across decompositions. As the range of the standard error of the means of each curve was small (pICA: 0 to 0.091, EEG: 0 to 0.0263, fMRI: 0.0012 to 0.093), to preserve visual clarity, error bars were not shown. The figures for the other source similarity thresholds (0.5, 0.6, 0.7) are shown in the supplementary material in Section C.2. The performance of pICA was greater than 0.8 only in the spatiotemporal domain. This was observed in the case when the pICA similarity threshold was 0.7 and the number of sources was 43. It was also observed when the pICA similarity threshold was 0.8 and the number of sources was 43, 45, or 50 or more.

4.3.2 Experiment 2: (p)ICA performance as a function of source noise dispersion

The performance of (p)ICA as a function of the source noise dispersion (0 to 0.2, step 0.01) is shown in Figure 4.12. The source similarity domain was spatiotemporal and the source similarity threshold was 0.8. Each column corresponds to a pICA similarity domain (temporal, spatial, spatiotemporal) and each row to a pICA similarity threshold (0.5 to 0.8, step 0.1). For each simulation, ICA was performed 30 times and the curves correspond to the mean across decompositions. As the range of the standard error of the means of each curve was small (pICA: 0.002 to 0.019, EEG: 0.005 to 0.014, fMRI: 0.003 to 0.011), to preserve visual clarity, error bars were not shown. The figures for the other source similarity thresholds (0.5, 0.6, 0.7) are shown in the supplementary material in Section C.3. The performance of pICA was greater than 0.8 only in a single case when the pICA similarity domain was

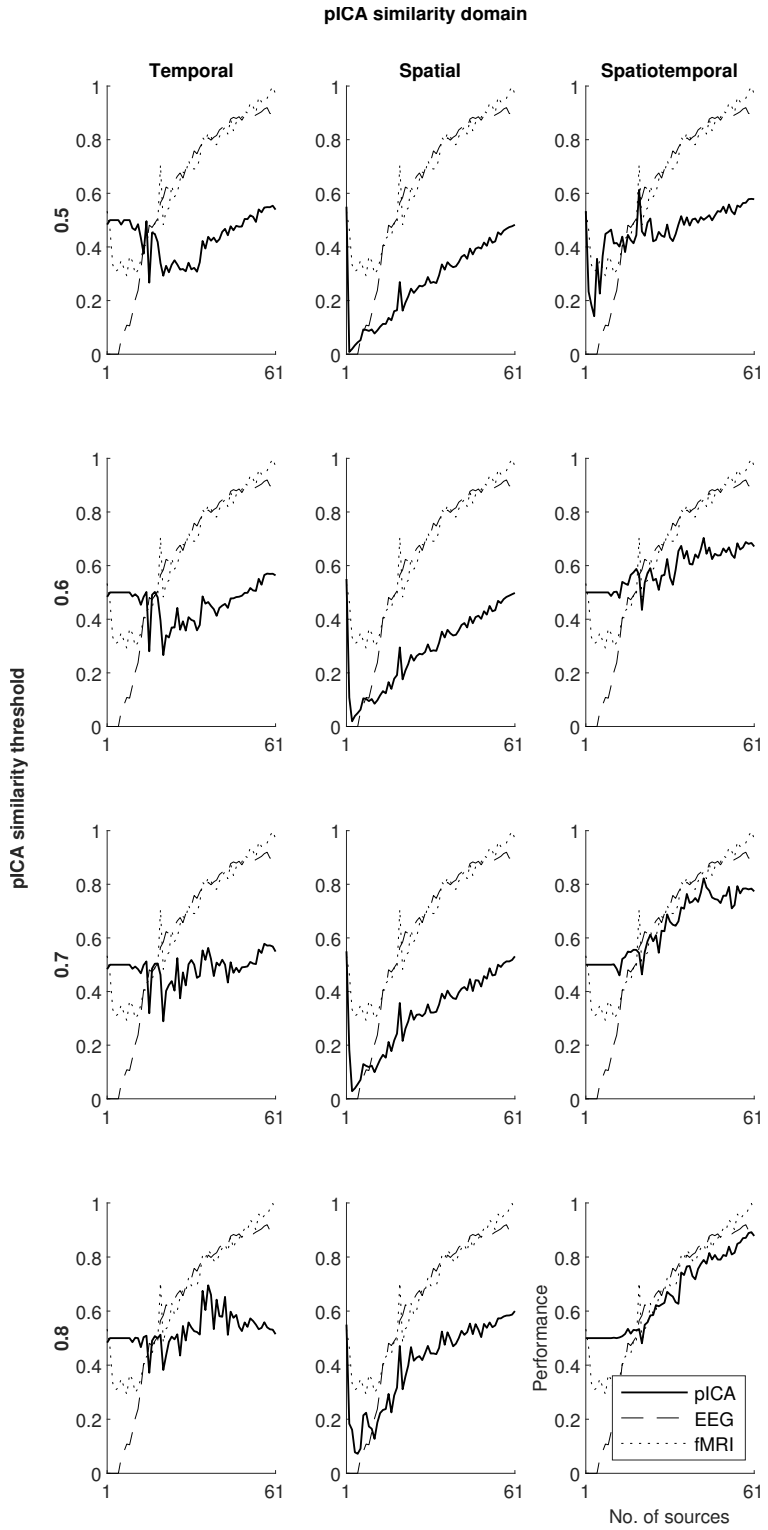

Figure 4.11 Experiment 1 results: (p)ICA performance as a function of the number of sources (1 to 61, step 1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.8. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0 to 0.091, EEG: 0 to 0.0263, fMRI: 0.0012 to 0.093) and are not shown to preserve visual clarity.

spatiotemporal, the pICA similarity threshold was 0.8, and there was no noise (dispersion = 0). In the presence of noise, the highest performance was 0.58, and this again was in the spatiotemporal domain with a threshold of 0.8 and the dispersion was 0.01.

4.3.3 Testing MI for measuring pICA similarity

The performances of EEG and fMRI ICA with correlation and MI as the source similarity measures across the source similarity threshold (0 to 1, step 0.1) are shown in Figure 4.13. For correlation, EEG ICA performance was high (> 0.8) when the threshold was between 0.5 and 0.9, and fMRI ICA performance was high when it was between 0.5 and 0.8. For MI, EEG ICA performance was high when the threshold was between 0.6 and 0.8, and fMRI ICA performance was high when the threshold was 0.4.

The performance of pICA with each of the measures as the source and pICA similarity measure across the source and pICA similarity thresholds is shown in Figure 4.14. The highest pICA performance (0.93) was observed when correlation was used for pICA similarity (threshold = 0.6) and MI for source similarity (threshold = 0.5). When both measures were correlation, the highest performance was 0.88 when the source similarity threshold was 0.6 and the pICA similarity threshold was 0.3. Generally, the performance was mostly greater than 0.7 when the source similarity threshold was 0.5 or 0.6, and when it was greater than 0.3 and the pICA similarity threshold was greater than 0.5. When both measures were MI, pICA performance was high (> 0.8) only when the source similarity threshold was 0.4. For other values of the threshold, the performance was greater than 0.7 when the source similarity threshold was 0.5 and the pICA similarity threshold was between 0.3 and 0.5.

4.4 Discussion

4.4.1 Experiment 1: (p)ICA performance as a function of the number of sources

The performance of pICA was high (> 0.8) in the spatiotemporal domain and mainly when the pICA similarity threshold was 0.8 and the number of sources was 50 or more. The most likely reason that pICA performed well at this threshold is that there were the least number of pICs and, therefore, the specificity of the pICs was high.

In previous literature, pICA has only been applied in the temporal domain, and the results show that it did not perform well (> 0.8) in any case in this domain. The highest performance was 0.695 when the pICA similarity threshold was 0.8 and the number of sources was 37,

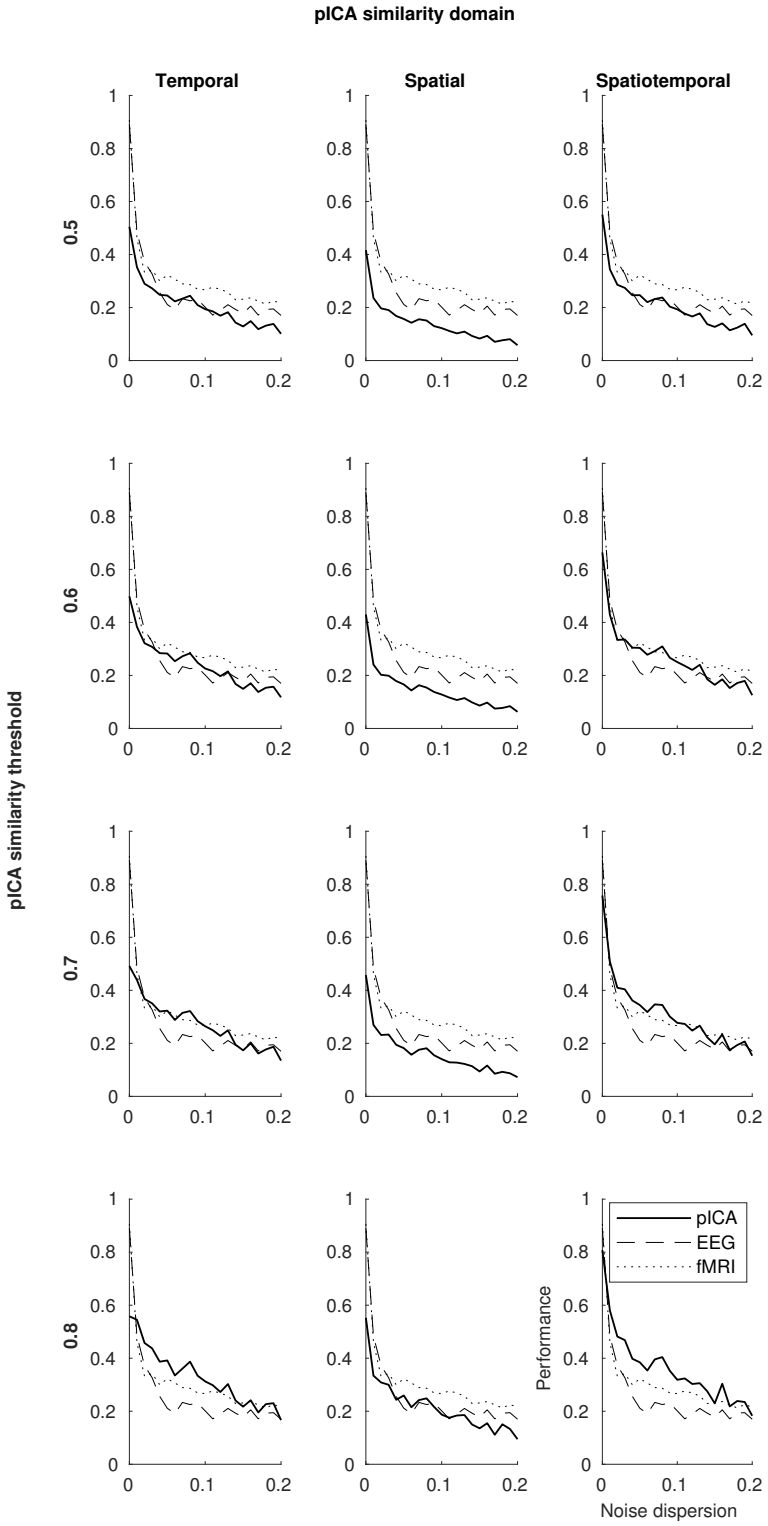

Figure 4.12 Experiment 2 results: (p)ICA performance as a function of the source noise dispersion (0 to 0.2, step 0.01) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.8. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0.002 to 0.019, EEG: 0.005 to 0.014, fMRI: 0.003 to 0.011) and are not shown to preserve visual clarity.

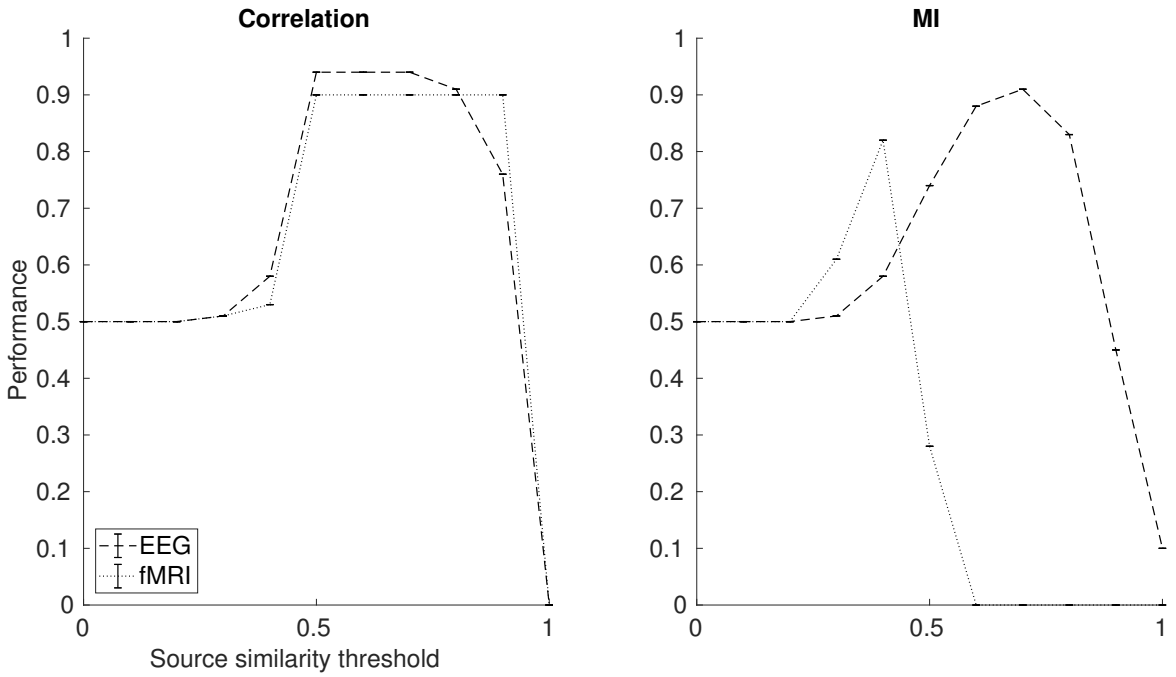

Figure 4.13 EEG and fMRI ICA performances with correlation (left) and MI (right) as the source similarity measures across the source similarity threshold (0 to 1, step 0.1) in the spatiotemporal source similarity domain. These were obtained from a single EEG ICA and fMRI ICA decomposition on signals from a simulation with 50 orthogonal sources with no source or sensor noise.

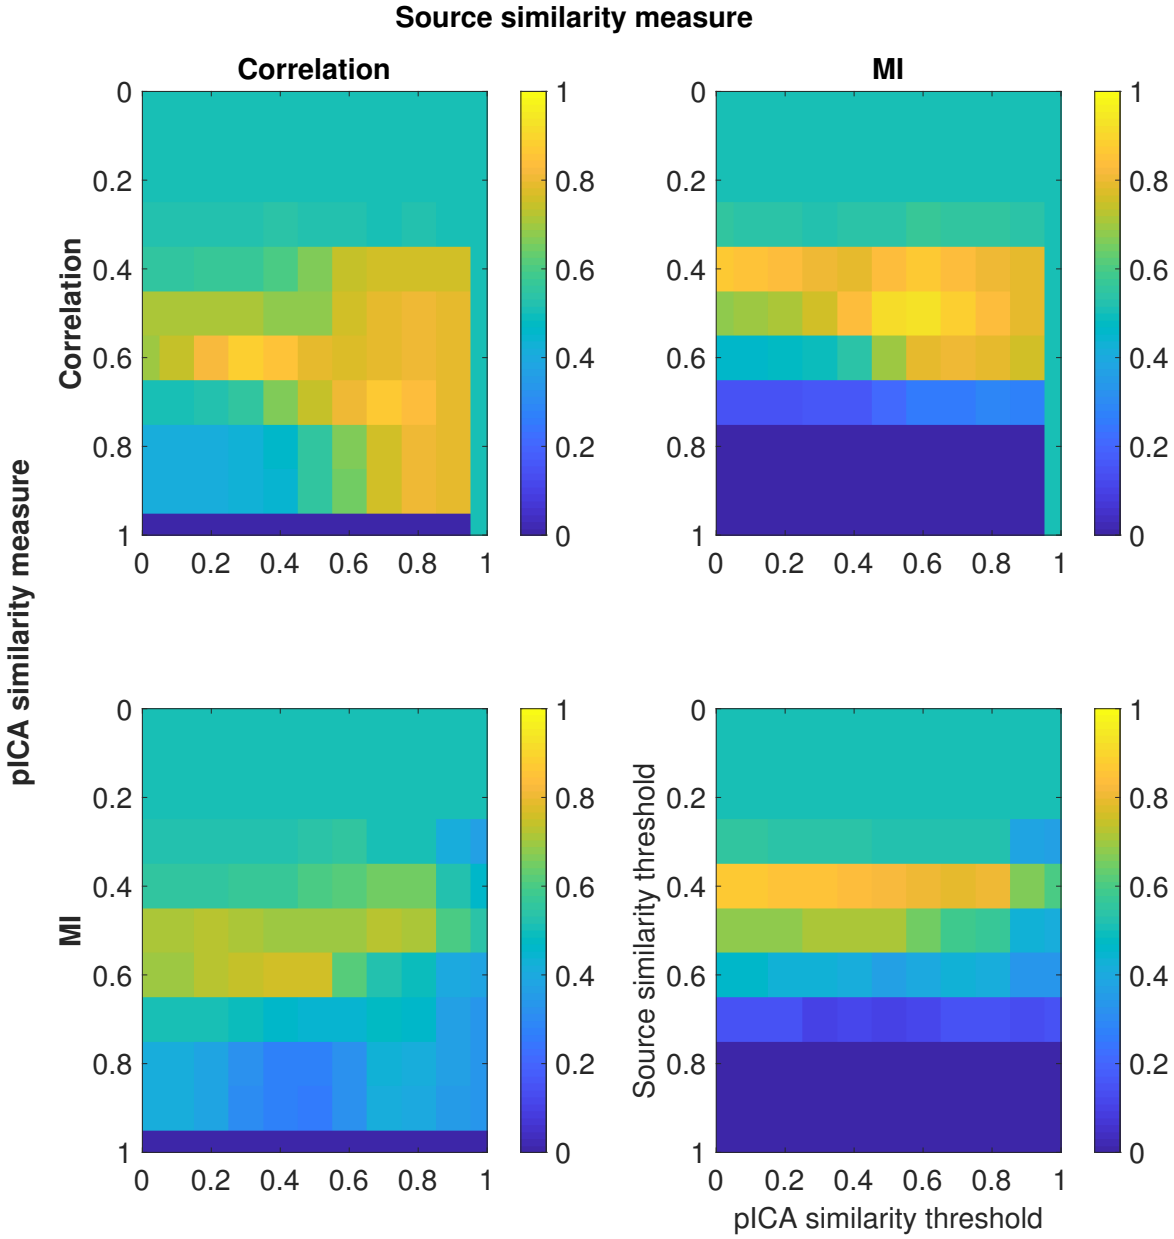

Figure 4.14 The performance of pICA with correlation (top row) and MI (bottom row) as the pICA similarity measures across pICA and source similarity thresholds (0 to 1, step 0.1) in the spatiotemporal pICA and source similarity domains. As it was unknown which measure more accurately mapped sources to ICs, both were used as the source similarity measures (left column: correlation, right column: MI). The performance of pICA was evaluated on a single EEG ICA and fMRI ICA decomposition on signals obtained from a simulation with 50 orthogonal sources with no source or sensor noise.

and, overall, the performance was > 0.6 when the pICA similarity threshold was 0.8 and the number of sources were around 40. It should be mentioned that, from a pragmatic point of view, one reason why pICA has not been performed in the spatial domain is that it is less straightforward to perform spatial matching between the EEG and fMRI ICs than it is to perform temporal matching, as, for the former, subject-specific projection matrices that map electrical activity within fMRI voxels to scalp locations are needed. However, solutions for estimating these using models of current propagation within the various kinds of tissue in the head (e.g. brain, skull, scalp) exist (e.g. Delorme and Makeig, 2004; Oostenveld et al., 2011) and could be used for this purpose.

To check if the reason for pICA performing well in the spatiotemporal domain was that this was also the source similarity domain, the experiment was repeated in the temporal (Figure C.5 in the supplementary material) and spatial (Figure C.6 in the supplementary material) source similarity domains with the source similarity threshold set to 0.8. It was observed that pICA performance exceeded 0.8 in some cases when the pICA and source similarity domains were temporal and when they were spatial, confirming that pICA performance was better when it was evaluated in the same domain it was applied in, suggesting a potential bias in the performance measure. However, pICA performance exceeded 0.8 in the spatiotemporal pICA similarity domain even when the source similarity domain was not spatiotemporal, and pICA performance trends in this domain were largely unchanged across the source similarity domains. Also, in the temporal source similarity domain, EEG ICA performance increased, whereas fMRI ICA performance did not, and the opposite was observed in the spatial domain, whereas in the spatiotemporal domain both performances increased. These observations were considered to be further evidence that choosing the spatiotemporal source similarity domain for evaluation had been the optimum decision.

Based on Chapter 2, it had been hypothesised that fMRI ICA performance would be high and unaffected by the number of sources and EEG ICA performance would linearly increase with the number of sources. It was observed, instead, that both EEG and fMRI performances linearly increased with the number of sources, with the exception being that fMRI ICA performance was fairly stable around 0.3 when the number of sources was between 3 and 13. To further understand these curves, the the two quantities used to calculate the performance measures (see Section 4.2.5) were plotted for the case where both domains were spatiotemporal and the thresholds were 0.8. (Figure 4.15). For EEG and fMRI ICA, the quantities were almost exactly the same, and, after inspection, this was confirmed to be the case because the matching sources and ICs mostly had a one-to-one mapping. While the blue curve in Figure 4.15 resembles the performance measure used in Chapter 2 in the sense that it describes the proportion of sources that have at least one IC with a correlation

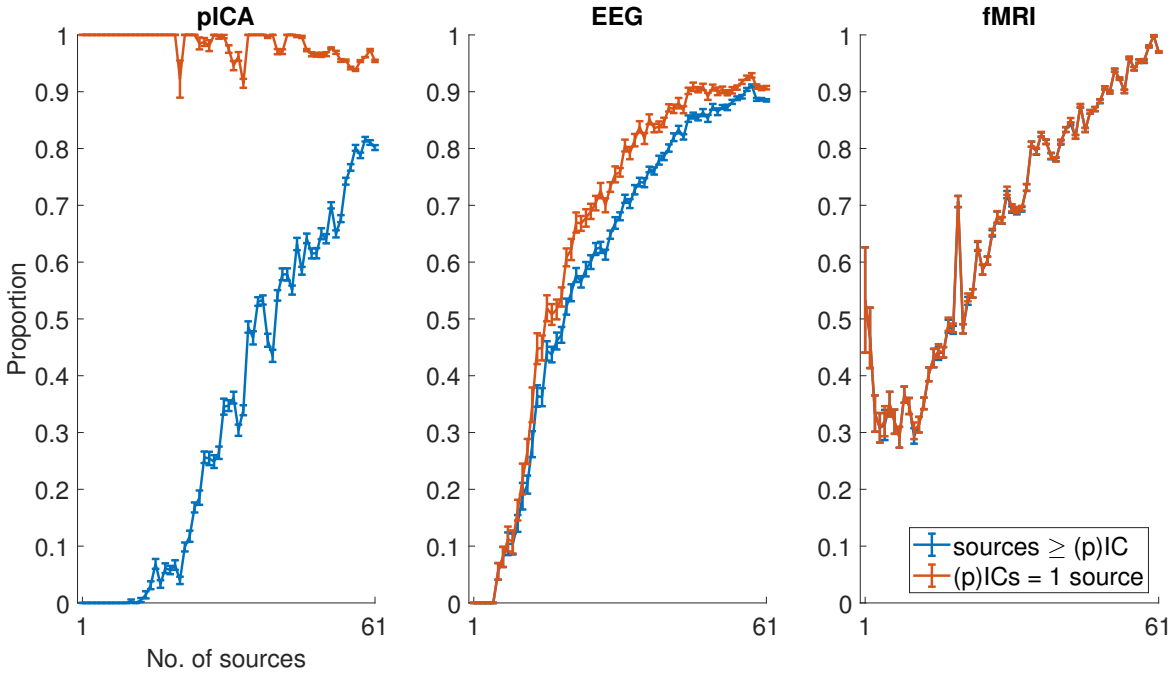

Figure 4.15 The proportion of sources with at least one matching (p)IC (blue curve) and the proportion of (p)ICs matching exactly one source (orange curve) for pICA, EEG ICA, and fMRI ICA. The pICA and source similarity domains are spatiotemporal and the thresholds are 0.8. The mean of these two quantities is used as the performance measure (see Section 4.2.5). These quantities were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions and the errors bars to the standard error of the means.

greater than the source similarity threshold (0.8 in this case), the trend of this curve for fMRI did not match the results observed previously either. The two main differences between the previous measure and the blue curve are that (i) the source similarity domain for the previous measure was temporal whereas for the current measure it is spatiotemporal, and (ii) the previous measure only looked at the performance of the first source, whereas the current one looks at the overall proportion of sources. To check if the former difference explained the inconsistency, the fMRI blue curve was inspected for the case where the source similarity threshold was temporal (Figure C.7 in the supplementary material) and in it was observed to still not resemble the previous results. Therefore, the latter difference is currently understood to provide the explanation.

Based on the EEG ICA hypothesis, it was expected that pICA performance would linearly increase with the number of sources. Figure 4.15 shows that the specificity of pICA was always high, and the linear trend observed in pICA performance was driven by the performance of ICA in each modality to extract ICs that corresponded to sources. This figure also illustrates that pICA performance can exceed both EEG and fMRI ICA performances as pICA has higher specificity than single-modality ICA, as for a pIC to match a source requires both the EEG IC and the fMRI IC to have high similarity with the source.

The results of this experiment suggest that pICA should be applied in the spatiotemporal domain, pICA performance increases with the number of sources, and for good performance, the number of sources should be around 45 or more. The pICA similarity threshold did not largely affect pICA performance, and the reason for this is could have been the orthogonal construction of the sources.

4.4.2 Experiment 2: (p)ICA performance as a function of source noise dispersion

This experiment had the fairly straightforward hypothesis that (p)ICA performance would decrease with source noise dispersion. The purpose of this experiment was to see how robust pICA was to noise in order to determine its viability as a method for combining joint EEG-fMRI signals. As hypothesised, (p)ICA performance decreased with noise dispersion, with pICA not performing well (> 0.8) when there was noise (dispersion > 0) in any case. When lowest non-zero value of dispersion was used (0.01, SNR = 0.51), the highest pICA performance observed was 0.58, EEG ICA performance was 0.484, and fMRI ICA performance was 0.46. These were all when the pICA similarity domain was spatiotemporal and the pICA similarity threshold was 0.8.

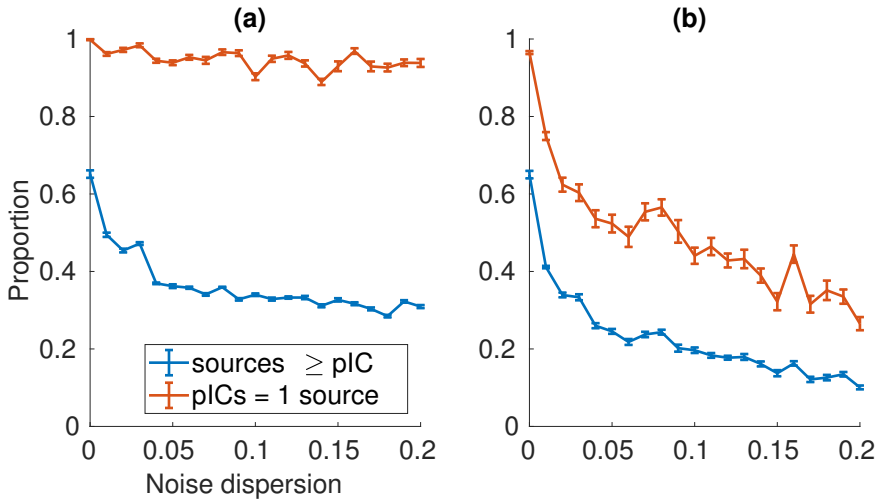

Figure 4.16 Comparison of the proportion of sources with at least one matching pIC (blue curve) and the proportion of pICs matching exactly one source (orange curve) when the source similarity threshold is (a) 0.7 and (b) 0.8. The pICA and source similarity domains are spatiotemporal and the pICA similarity threshold is 0.8. These quantities were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions and the errors bars to the standard error of the means.

For reasons described in Section 4.2.6, the source similarity threshold was set to 0.8. It was observed, however, that (p)ICA performance was significantly higher at lower thresholds (Figures C.8, C.9, C.10 in the supplementary material). To better understand this, the two quantities used to calculate pICA performance were plotted in Figure 4.16 when the source similarity thresholds were (a) 0.7 and (b) 0.8. From this figure, it was identified that the main difference between the two cases was that the pICs had high specificity when the source similarity threshold was 0.7 compared to when it was 0.8. As increasing the source similarity threshold will reduce the number of sources that match a pIC, the drop in pIC specificity with the increase in threshold suggests that a significant proportion of the pICs stopped matching any source as a result of the increase in threshold. A qualitative evaluation comparing some of the pICs for the two thresholds with their corresponding sources was not performed. This analysis will be useful to determine if a lower source similarity threshold is more appropriate. In Chapter 2, it had been observed that EEG ICA had a performance around 0.55 and fMRI ICA had a performance around 0.7 when the noise dispersion had been 0.2 (Experiment 4 in Chapter 2). In the current work, both of these were around 0.2 (EEG was 0.17, fMRI was 0.22). One explanation for the inconsistency may be that in Chapter 2 the performance of the first source was only evaluated and this source had a size of 37 regions, whereas in this work the performance measure accounted for all 50 sources and all of them comprised single regions.

The results of this experiment suggest, similar to the previous one, that the spatiotemporal domain is the most suitable for performing pICA and the similarity threshold does not largely affect performance. Regarding the robustness of pICA to source noise, pICA performance drops significantly in the presence of noise when the source similarity threshold is 0.8, but remains relatively high (> 0.6) even when the dispersion is 0.2 (SNR = 0.08) when the source similarity threshold is dropped to 0.7. Therefore, further investigation should be performed to check if a lower source similarity threshold would be more appropriate.

4.4.3 First look at using MI for pICA

The examination of MI for pICA was exploratory and there was not any hypothesis on whether MI would perform better than correlation or not, and the motivation simply was to check if the measure provided different results than correlation, and, if it did, could these results be considered to be more accurate? The EEG and fMRI signals from a single simulation with 50 orthogonal sources and no source or sensor noise were used, and the ICs from single EEG ICA decomposition and fMRI ICA decomposition were analysed.

The pICA similarity matrices of both measures were fairly similar and there were cases where one measure had high values of similarity for a pair of ICs for which the other had low. Two such cases were inspected and it was qualitatively judged that in both cases correlation was more accurate than MI.

EEG and fMRI ICA performances with the two measures were evaluated. With correlation, as had already been noted in Section 4.2.6, both modalities mostly performed well (> 0.8) when the source similarity threshold was greater than 0.5. With MI, while EEG ICA performance was relatively similar to its correlation performance, fMRI ICA performance was significantly lower, with it only being higher than 0.8 in a single case, when the threshold was 0.4. To understand this further, the two quantities used to calculate the performance measures (see Section 4.2.5) were plotted (Figure 4.17). This figure showed that when the source similarity threshold was 0.4, there was a peak in the number of ICs that matched exactly one source. It also showed that there was a sharp decrease in the number of sources having at least one matching IC. Both of these processes together explain why fMRI ICA performance peaks when the threshold is 0.4 and sharply decreases afterwards. This trend in fMRI ICA performance also explains why, when MI was the source similarity measure, pICA performance was only high (> 0.8) around this threshold.

While it had been assumed that the most ‘fair’ comparison of the measures for pICA performance would be comparing the cases when they were both the pICA similarity measure and the source similarity measure, i.e. the top-left matrix for correlation and the bottom-right for MI in Figure 4.14, the highest pICA performance was observed when correlation was

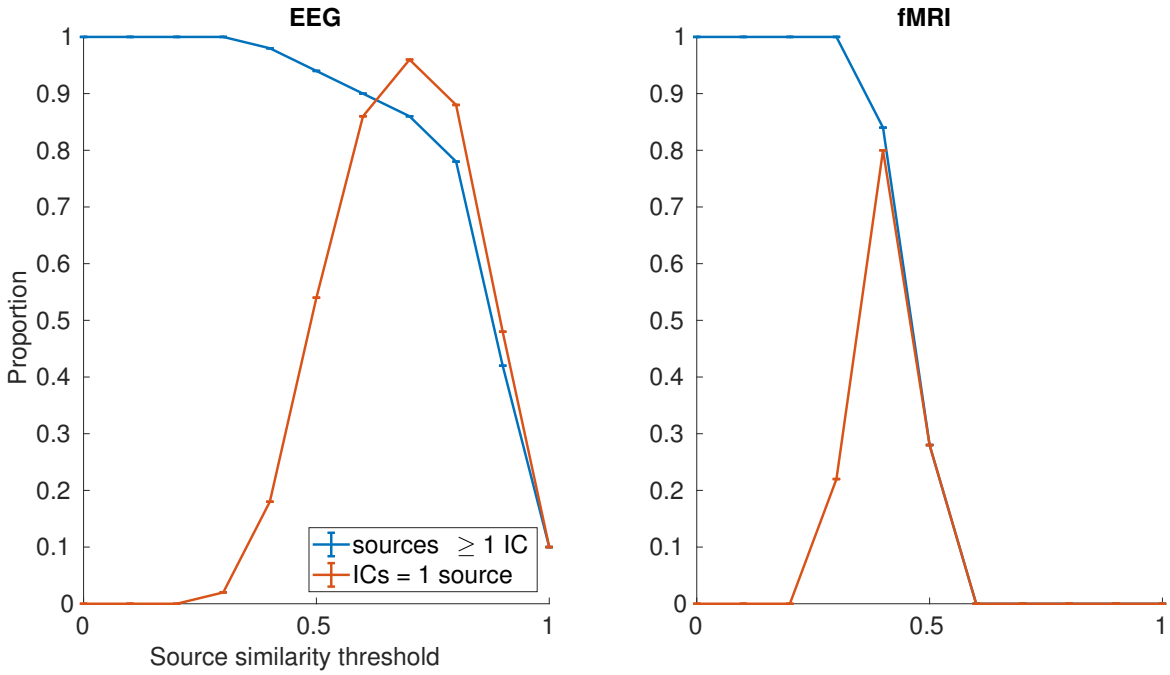

Figure 4.17 The proportion of sources with at least one matching IC (blue curve) and the proportion of ICs matching exactly one source (orange curve) for EEG ICA (left) and fMRI ICA (right) across the source similarity threshold (0 to 1, step 0.1) with MI as the source similarity measure. The source similarity domain is spatiotemporal. These were obtained for a single EEG ICA and fMRI ICA decomposition on signals from a simulation with 50 orthogonal sources with no source or sensor noise.

used for pICA similarity and MI for source similarity. The reasons behind have not been investigated further, but they highlight the dependence of pICA evaluation on the assumptions made for establishing the ground-truth source-to-IC mappings.

Overall, pICA had higher performance with correlation as the pICA similarity measure, irrespective of which measure was used for evaluating source similarity. In terms of accuracy, correlation was also the 'superior' measure, based on the two cases that were inspected where the measures were inconsistent in their values of similarity between EEG-fMRI IC pairs. Further, more systematic, investigation should be performed to validate the generalisability of these results. The aim of exploring the application of MI for pICA in the current work was to provide a first look at its performance and to stimulate further research in this direction.

4.4.4 Limitations and future work

As highlighted in our previous work (Chapter 2), there are some limitations to the way stimuli have been constructed. One of them is that they are pulse-trains with three pulses each and, therefore, they model experimental paradigms where the activity of each source corresponds to a condition with only three trials, which is too many conditions with too few trials to be used in practice. The main reasons for this were that EEG ICA was observed to not perform well when the number of sources was low, and adding more pulses to the stimuli would increase the length of the simulated signals resulting in longer computational times and storage requirements. Another limitation is that the sources only comprise single, non-overlapping, regions and are temporally orthogonal, which is biologically implausible. This was done to characterise pICA performance in 'best-case' conditions in the first instance of this investigation, with the sources being highly separable based on their mutual independence. As a next step, further investigation with non-orthogonal sources should be performed, as this will be more biologically plausible, and it might also shed more light on how the pICA similarity threshold should be selected, which was not observed to largely affect pICA performance and this could have been due to the orthogonal construction of the sources.

Apart from the limitations pertaining to the stimulus parameters, several assumptions regarding the simulation parameters were made as well, e.g. the default TVB parameter values and connectome were used (see Chapter 2 for details), and all vertices had the same level of white Gaussian noise added. These assumptions were made as they were deemed to provide a reasonable starting point for this investigation. Further work with more diverse simulation parameters will be useful for inferring the generalisability of the results.

One cause for some concern is that the MI values that were observed were very low, typically less than a bit, and it does not seem plausible that the EEG and fMRI ICs shared such a small amount of information. The primary reasons for choosing MutualInfo 0.9 (Peng

et al., 2005) to calculate these values were that its accompanying article is highly cited and it is straightforward to use. While a comparison of MI toolboxes is beyond the scope of this work, based on the range of the MI values observed, we think that it will be useful to replicate the work described in Section 4.2.9 using a different toolbox and comparing the results. As all MI implementations rely on heuristics for estimating probability density functions, a computationally feasible method for calculating MI with complete accuracy does not exist, and, therefore, the results presented in this work with MutualInfo 0.9 are appropriate for a first inspection of pICA performance with MI as the pICA similarity measure.

4.5 Conclusions

This work is the first evaluation of 'trial-by-trial' pICA using simulated EEG-fMRI signals. The motivation was to characterise how pICA, in its most basic form, performed, and to test if its performance improved with the incorporation of spatial information and MI, neither of which have previously been used. It was also to demonstrate how TVB, which is a comprehensive human brain simulator that operates at multiple temporal and spatial scales and implements local and long-range connectivity, can be used for the evaluation of multimodal methods. The performance of pICA was observed to increase with the number of orthogonal sources in the absence of noise, and to decrease with the level of noise dispersion when the number of sources was fixed. The incorporation of spatial information improved pICA performance, which was generally higher when the EEG and fMRI ICs were matched in the spatiotemporal domain rather than just in the temporal domain. MI, however, was not observed to be a better measure than correlation in matching the ICs. This work concludes that pICA can provide accurate source information, especially when the number of sources is high (around 45 or more) and the level of source noise dispersion is low, when the sources are orthogonal. For further characterisation of pICA performance, future work should address the questions raised, systematically increase the mutual dependence between the sources, and increase the diversity of the simulation parameters.

Chapter 5

General Discussion and Conclusions

Concurrent EEG-fMRI is currently one of the most optimal means for non-invasively acquiring measurements of human brain activity with high spatial and temporal resolution. Parallel ICA (pICA) is a framework for analysing these signals that combines information from the two modalities by first performing ICA on each modality to obtain ICs that provide modality-specific source information and then matching the ICs across modalities based on their statistical similarities. Together, the matched ICs provide multimodal source information, with the EEG IC providing a high resolution temporal description and the fMRI IC providing a high resolution spatial description. In this thesis, EEG ICA, fMRI ICA, and 'trial-by-trial' EEG-fMRI pICA (see Section 1.5.1.2), collectively referred to as (p)ICA, were evaluated in terms of their abilities to provide source information about underlying functional networks in the brain. Two novel extensions to pICA, which are matching the ICs across modalities using spatial features and mutual information (MI), were also developed and evaluated.

This project started out with the broad aim to develop a method for analysing concurrent EEG-fMRI signals. After a review of the landscape of EEG-fMRI methods, it was decided to work on extending pICA, as it uses a simple, model-free[1] approach and there are several ways in which it could be further developed. It was also decided that before we could develop pICA, we needed to construct a framework for measuring its accuracy so that we could evaluate the assumption of pICA that some sources have corresponding ICs in both modalities and evaluate our extensions to pICA to see if they provided improvements. Therefore, with this in mind, the initial roadmap was: (i) obtain an EEG-fMRI dataset for which ground-truth source information is known that can be used for the evaluation, (ii) evaluate ICA in each modality to check the above-mentioned assumption, (iii) evaluate the most basic, 'stripped-down'

[1] ICA is not completely model-free as it assumes linear mixing and that at most one source is Gaussian (see Section 1.4).

version of pICA, and (iv) systematically extend pICA by evaluating its performance after each increment. During the course of this project, as described in more detail in the next section, this roadmap was mostly followed: The Virtual Brain (TVB) (Sanz-Leon et al., 2015b, 2013) was used to produce multiple, synthetic datasets for (i); (ii) was performed in Chapter 2; and (iii) and (iv) were performed in Chapter 4, with the only two extensions that were developed and evaluated being the addition of spatial features and mutual information.

In this chapter, a summary of the thesis is provided in Section 5.1. The observation that ICA and pICA performances increased with the number of sources is discussed in some detail in Section 5.2, with the acknowledgement that the reasons behind this are still not understood. A reason why the fMRI performance measure used in Chapter 2 might not be completely reliable is described in Section 5.3. The limitations of this project and corresponding future investigations are described in Section 5.4. The novel contributions of this project are stated in Section 5.5, and a summary of the conclusions and the take-home messages are described in Section 5.6.

5.1 Thesis summary

5.1.1 Evaluation of ICA

In Chapter 2, the performance of ICA in each modality was evaluated. The motivation behind this was to evaluate the assumption of pICA that some sources have corresponding ICs in both modalities and to obtain a dataset that was known to perform well with ICA in both modalities so that it could be used to evaluate pICA. As work on this progressed, it was decided that the evaluation of ICA in each modality was a useful contribution on its own and, therefore, this chapter was prepared as an independent manuscript. In the process, we realised that rather than evaluating pICA on a single dataset on which single-modality ICA was known to perform well, it would be more useful and generalisable to demonstrate its performance on a range of datasets that systematically varied in their parameter values. Another motivation for evaluating single-modality ICA was to address the more fundamental question on how (p)ICs should be interpreted - while EEG and fMRI ICs are typically interpreted as providing 'source' information, what does is this source information relate to in biological terms? To answer this question, we first defined sources as networks of regions with instantaneously synchronised local field potentials (LFPs), and we then evaluated ICA in each modality in terms of the extent to which the ICs provided information about these sources. Put together, we addressed the question if ICs could be interpreted as corresponding to independent source networks of regions with instantaneously synchronised LFP activity.

The parameter space we could explore was very large. Therefore, to keep the number of variables small, default values were used for all the intrinsic TVB parameters (see Table 2.1), such as those pertaining to the neural mass models and connectivity. The only extrinsic parameters that were varied pertained to the construction of the sources in terms of their number, time-courses, spatial maps, and noise. We chose to characterise ICA performance in terms of the number of sources, the size of a source, and the level of source noise, as these parameters were deemed to be the most suitable to start the investigation with. In this chapter, it was observed that: EEG ICA performance improved with the number of sources, whereas fMRI ICA performance was relatively unchanged; ICA performance of both modalities decreased with the level of noise, and varied without clear trends with source network size in the presence of noise. One of the key take-home messages of this chapter was that there are cases when ICA performance, especially EEG ICA performance, is low in ideal conditions, i.e. when the sources are orthogonal and there is no noise and, therefore, ICs should be interpreted cautiously.

5.1.2 Inspecting the generalisability of the source-to-region mappings

After Chapter 2 was completed, further investigations were performed to assess the impact of the source-to-region mappings on ICA performance (see Appendix A). It had been assumed that ICA performance would be largely unaffected by the anatomical locations of the sources as long as their contributions could be detected in the EEG and fMRI signals. The reason for this was that ICA is blind to the underlying forward models that map the source activities to the EEG and fMRI signals and separates the sources from the mixtures only based on their independence. However, two observations prompted further investigation. The first was that EEG ICA performances in Experiments 1 and 2 were inconsistent when the first source comprised a single region even though the number of sources in both experiments was 50 (see Section A.2). The second was the dip observed in fMRI ICA performance in Experiment 1 (see Figure 2.8), which seemed to be an anomaly (see Section A.3).

To check if the anatomical location of the first source affected ICA performance, an experiment was run in which the source comprised each region in the TVB parcellation (Sanz-Leon et al., 2015a) individually across simulations (see Section A.2). It was concluded from this experiment that EEG ICA performance varied not only depending on the region the source comprised but also depending on the regions the other sources comprised as well. To check if the results of Experiment 1 contained the dip in fMRI ICA performance when the source-to-region mappings were varied for all sources, it was repeated an additional nine times with different source-to-region mappings, and it was observed that in most cases the dip was visible and the mean trend across the replications was similar to what had been observed

originally in the experiment (see Section A.3). These investigations raise further questions as they are somewhat contradictory. The first investigation showed that the source-to-region mappings did affect EEG ICA performance, whereas the second investigation showed that the source-to-region mappings did not largely affect ICA performance.

To better understand this and to check how robust the results of Experiment 1 are (Figure 2.8) using the measure from Chapter 2, the experiment in Section A.3 should be repeated for each region of the first source with multiple source-to-region mappings for the other 'non-target'[2] sources. In the evaluation of pICA (Chapter 4), while the same simulation design was used, the performance measure did not only account for the performance of the first source, but rather it accounted for the performance of all sources. For this reason, the new performance measure is considered to be more robust to the source-to-region mappings. However, this should be examined in future investigations.

5.1.3 Evaluation and development of pICA

As described earlier, the goal of this project was to develop pICA and it was planned to do this systematically by first identifying the most basic, 'stripped-down' version of the method, evaluating its performance, and then incrementally adding more features and checking if they improved its performance. The 'stripped-down' version of the algorithm was identified to be the one that worked on the data of a single subject and that performed no temporal reduction on the EEG or fMRI signals or on the IC time-courses. It was considered important to evaluate pICA on the data of a single-subject as it demonstrates how sensitive it is to within-subject effects without the additional complexity group ICA brings by concatenating the data from multiple subjects, performing dimensionality reduction across subjects, and matching the ICs across modalities using their back-projections (see Appendix D). Initially, there were several features that we had planned to add to pICA. These included: matching in the spatial domain; matching using a variety of connectivity measures (see Appendix E); and comparing single-subject pICA with pICA that used various configurations of group ICA (see Erhardt et al. (2011) for review). As the work in Chapter 2 took considerably longer than anticipated, the extensions to pICA were limited to performing matching in the spatial domain and using mutual information.

In Chapter 4, pICA performance was evaluated as a function of the number of sources and the level of source noise. In Chapter 2, ICA performance had been evaluated as a function of source network size as well (Experiments 2, 5, and 6). This was not performed in Chapter 4 for two reasons. One was that we decided to start with a smaller set of experiments for

[2]Sources for which ICA performance is not evaluated, i.e. all the sources except the first one.

our first investigation on pICA performance for the sake of simplicity and computational tractability. The other reason was that the performance measure used in Chapter 4 evaluated (p)ICA (i.e. EEG ICA, fMRI ICA, and pICA) performance of all sources as a whole, rather than just a single source and, therefore, it was not suitable to evaluate (p)ICA performance as the size of the first source varied. That said, this measure could have been modified for this purpose, by setting the value of the first quantity (see Section 4.2.5) to 1 if any (p)IC matched the first source and to 0 otherwise. However, it was decided not to do this in the first instance of this investigation in order to keep the measure consistent across experiments.

The performances of EEG and fMRI ICA observed in Chapters 2 and 4 and the performance of stripped-down pICA in Chapter 4 were consistent in that (p)ICA performance was observed to increase with the number of orthogonal sources in the absence of noise, and to decrease with the level of noise when the number of sources was fixed. The incorporation of spatial information to match the ICs across modalities, which was a novel feature, improved pICA performance, whereas the use of mutual information did not, and pICA performance was in fact significantly worse when mutual information was used instead of correlation.

5.1.4 Evaluation of ICA in the presence of source noise

Chapter 3 was a slight digression from the original aims of this work (see Section 1.7), but it proved to be a useful exercise. In this chapter, it was realised that a limitation of the performance measure used in Chapter 2 (and subsequently Chapter 4) was that it did not account for how much source noise an IC contained. This is only a limitation if the source noise is considered to be of interest, which was not the case in Chapters 2 and 4, as the focus of both of these chapters was on stimulus-evoked activities. However, to investigate how ICA could be evaluated while accounting for noise if this were of interest, two approaches were explored. The first approach was very similar to the existing approach used in Chapter 2 with the difference being that the time-course of the first source was not approximated by its stimulus time-course but rather by the mean LFP signal across the vertices it comprised. However, in terms of the source time-course being explained by a linear combination of the EEG signals, this approach was only marginally better than the one used in Chapter 2 (see Section 3.2.3), indicating that the mean LFP signals did not account for source noise significantly better than the stimulus time-courses and, therefore, would not have been a significant improvement over the stimulus time-courses for evaluating if the ICs contained the source noise in addition to the stimulus-evoked activities. The second approach was significantly different from the previous methods. In this approach, for the source, a predicted IC time-course was constructed using its stimulus time-course and the source noise time-course at each vertex it comprised. This method, in some cases, had much

higher ICA performance than the method used in Chapter 2 (Figure 3.18). However, while promising, this method was not explored further as it was computationally infeasible to run on a large-scale and because, in this work, the evaluation of the ICs in terms of how much source noise information they provided was not of primary interest, as the focus was on task-evoked (i.e. stimulus-evoked) brain activity. While the work performed in this chapter did not have practical consequences on the work described in Chapters 2 and 4, it was useful as it: (i) made us explicitly aware that our evaluation of (p)ICA was limited to only measuring how much information the ICs contained about the sources in terms of their stimulus-evoked activity and not their source noise, and (ii) it provided an initial investigation into how ICA could be evaluated if the source noise were of interest.

5.2 Why does (p)ICA performance increase with the number of sources?

In Chapter 2, EEG ICA performance was observed to increase with the number of sources (Figure 2.8), and in Chapter 4, EEG ICA, fMRI ICA, and pICA performances increased with the number of sources (Figure 4.11). The reasons for the increase in ICA performance with the number of sources is unclear and, to our knowledge, there is not much literature that systematically investigates ICA performance as a function of the number of sources. One example in which this has been done is Giannakopoulos et al. (1999) where, contrary to what we observed, ICA performance, measured using the Amari index (Amari et al., 1996), decreased as the number of sources increased from 1 to 14. In this paper, the possible reasons for this were not discussed.

In Chapter 2, one reason for the increase in EEG ICA performance could be that as the number of sources increased, so did the number of ICs, as they were both set to be equal. As the number of ICs increased, the probability of there being a single IC that matched the first source might have increased as well. However, this does not explain the increase in performance seen for both modalities in Chapter 4 as, in that case, the performance measure accounted for the performance of all sources and not just a single source.

One interesting observation is that ICA performance increases with the number of sources in the source similarity domain in which the modality has higher resolution. For example, only EEG ICA performance increased with the number of sources in Chapter 2 and the source similarity domain was temporal and, in Chapter 4, only EEG ICA performance increased when the source similarity domain was temporal (Figure C.5) and only fMRI ICA performance increased when the source similarity domain was spatial (Figure C.6). At this

stage, it is unclear to us why this is case, but intuitively we feel that this is an important observation.

5.3 Reliability of fMRI ICA performance reported in Chapter 2

In Chapter 2, it was observed that fMRI ICA performance was nearly always higher than that of EEG ICA and, in retrospect, it was considered that one reason for this could have been that the fMRI ICs had lower specificity than the EEG ICs as the similarity between the ICs and the sources was measured in the temporal domain. To recap, in the temporal domain, the correlations between the predicted BOLD signals of the sources and the fMRI IC time-courses are measured. As the predicted BOLD signals of sources with small differences in their phase offsets (see Section 2.2.2.1) would have been highly correlated, due to the large overlap in their haemodynamic responses, it was considered possible that the first source could have been reported to have had high performance even when its corresponding IC did not exist in the decomposition as it could have had a high correlation with an IC corresponding to a source that had pulses close in time to it. As described in this section, it was concluded after further analysis that the fMRI ICA performance reported in this chapter could be considered to be reliable.

To illustrate the problem, the predicted BOLD signals of the first five sources are shown in Figure 5.1a for a simulation with 50 sources and the correlations between all the sources are shown in Figure 5.1b. To identify the threshold at which fMRI ICA performance can be trusted to indicate that the IC corresponded to the first source, the correlation between the predicted BOLD signal of the first source and all the other sources was calculated. Then, the number of sources that had a correlation with the first source that was greater than or equal to a threshold (0 to 1, step 0.01) was counted. This number indicates what the maximum number of ICs could be that correspond to a source other than the first one when fMRI ICA performance is equal to the threshold. This is shown in Figure 5.1c. As can be seen from the figure, none of the sources had a correlation with the first source that was greater or equal to 0.84, indicating that the fMRI performance measure could be trusted to indicate that the IC definitely did correspond to the first source only when its value was greater than or equal to 0.84. In Chapter 2, fMRI ICA performance was lower than this threshold mainly in Experiments 4 and 5 (Figures 2.11 and 2.12a respectively), with a small number of cases in the other experiments. In terms of interpretation, what this initially suggested was that when fMRI ICA performance was below 0.84, it was possible that the ‘actual’ fMRI ICA

performance was lower and, therefore, in these cases, the fMRI ICA performance measure should not have been considered to be mostly reliable.

However, further analysis revealed that the probability of the fMRI ICA performance value being unreliable was low. For this analysis, the proportion of ICA decompositions (out of 30) was calculated where the first source matched the same IC in the temporal and spatial domains for each simulation used for Experiment 2 in Chapter 4 (see Section 4.2.8). This is shown in Figure 5.2. As can be seen from the figure, the same IC matched the first source in both the temporal and spatial domains in most cases. Across the simulations, the mean proportion of matches across simulations was 0.93 and the standard deviation was 0.057. This indicates that the probability of the fMRI ICA performance value corresponding to an IC which the first source did not match to is very low as, in most cases, the IC matched the first source in the spatial domain as well. Therefore, while it is possible that in some cases the fMRI ICA performance reported in Chapter 2 was based on the correlation between the first source and an IC that did not spatially correspond to it, the probability of this is low and, therefore, the fMRI ICA performance reported in the chapter can be considered to reliable.

There were a few reasons why the evaluation of fMRI ICA was performed in the temporal domain and not the spatial domain in Chapter 2 even though the spatial resolution of fMRI signals is much higher than their temporal resolution. One of these was that source spatial maps are binary variables (see Section 2.2.2.2) and Pearson's correlation is not typically used with binary variables. Also, when looking at source network size, the ground-truth spatial map would vary from a vector containing mostly 0s, when it comprised a single region, to a vector containing all 1s, when it comprised all the regions. Therefore, the variance of the vector would be in the shape of an 'inverted u', with the variance being very low when the size is small and gradually increasing until the spatial map contains an equal number of 0s and 1s, and then decreasing again till it completely contains 1s, at which point the variance would be 0 and the correlation between the source spatial map and the sIC spatial map undefined. It was considered that measuring ICA performance as a function of size using correlation was more appropriate in the temporal domain as the systematic change in the variance of the spatial maps across size may 'bias' the measure in the spatial domain. Another reason for using the temporal domain was that for real fMRI signals it is standard practice to identify task-dependent spatial ICs by matching IC time-courses and predicted BOLD time-courses, which are constructed from stimulus time-courses (e.g. Calhoun et al., 2008). Nevertheless, in Chapter 4, the fMRI signals were matched in the spatial domain as well in order to explore how evaluation in the spatial domain would work, and the observations reported in the chapter do not appear to suggest that this was inappropriate.

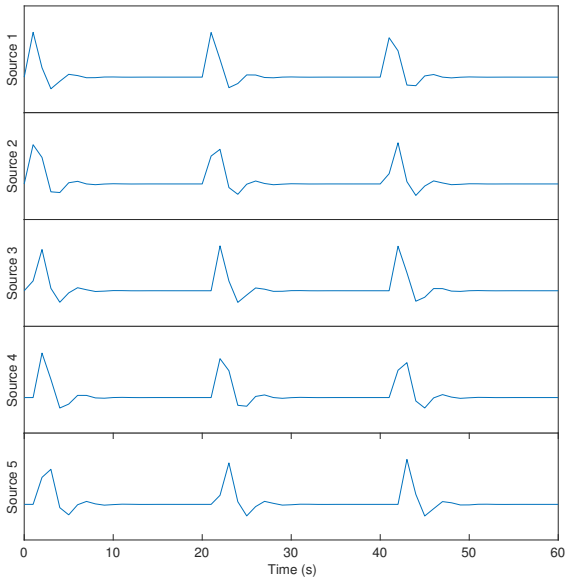

(a)

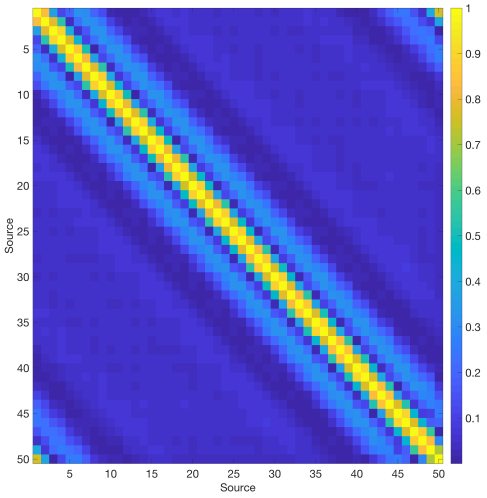

(b)

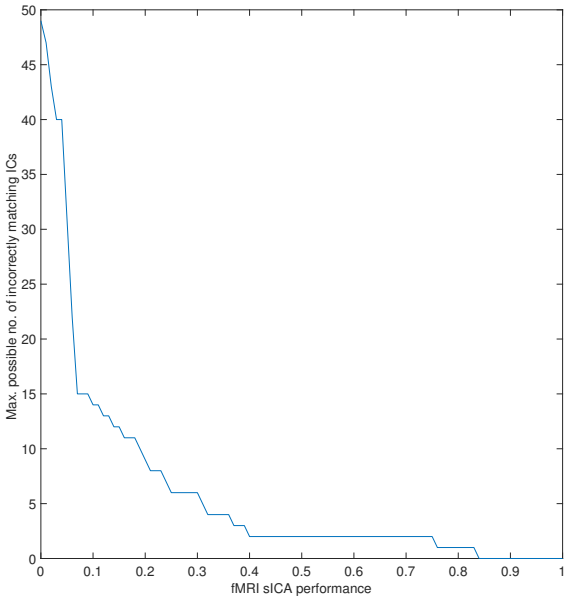

(c)

Figure 5.1 (a) The predicted BOLD signals of the first five sources in a simulation with 50 sources. (b) Correlations between the predicted BOLD signals of the 50 sources. (c) The possible number of ICs incorrectly matching the first source for each value of the fMRI ICA performance measure (0 to 1, step 0.01) used in Chapter 2.

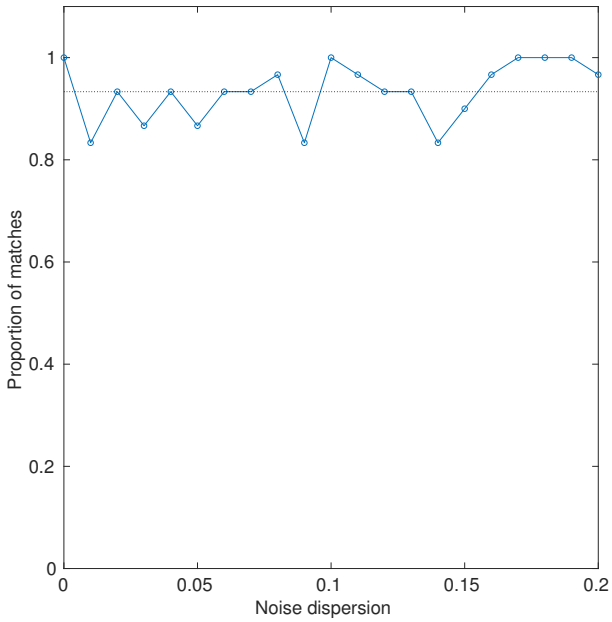

Figure 5.2 The proportion of ICA decompositions (out of 30) in which the first source matched the same IC in the temporal and spatial domains. The simulations varied in their noise dispersion (0 to 0.2, step 0.01, see Section 4.2.8). The mean proportion of matches across the simulations was 0.93 (horizontal line) and the standard deviation was 0.057.

5.4 Limitations and future work

5.4.1 Evaluation of (p)ICA with real data

In this project, (p)ICA was evaluated only using simulated data, and not using real EEG and fMRI signals. The main reason for this was that it is difficult to obtain ground-truth source information from real data without using invasive tools, e.g. intracranial EEG. One non-invasive experimental paradigm, however, that has been used in the validation of methods using real signals is retinotopic mapping (e.g Calhoun et al., 2001c; Capilla et al., 2016; Yoshioka et al., 2008). In this project, we collected EEG-fMRI data using this paradigm from five participants with the intention to use it for evaluating (p)ICA. We were unable to complete this part of the project as the evaluation with simulated signals took much longer than had been anticipated. The basic idea was to first obtain retinotopic mappings between checkerboard wedge positions and EEG and fMRI features and then to treat them as the ground-truth when evaluating (p)ICA. The exact details of how (p)ICA would be evaluated were not worked out before data was collected. The reason for this was that the data collection had to be expedited as the MRI scanner in the facility we use was about to be upgraded and we decided to collect the data before the upgrade as the timelines for when it would be

completed were not clear. A limitation of the work presented in this project therefore is that (p)ICA has not been evaluated using real data. However, in the future, we will attempt to evaluate it using this retinotopy dataset, provided that the dataset can be used for this purpose and we have sufficient resource.

5.4.2 Simulation parameters

As described in Sections 2.4 and 4.4.4, the generalisability of the results presented in Chapters 2 and 4 is limited as the intrinsic parameters (pertaining to simulator function) were fixed across simulations and only a small number of extrinsic parameters (pertaining to stimulus properties) were varied. These limitations are briefly summarised here again.

One of the limitations was that the sources only had pulse-train time-courses. These were used to model typical event-related experimental paradigms with 'on' encoding the presentation of a stimulus and 'off' encoding its absence. To increase the generalisability of the results to other types of sources, further work should be performed with more complex stimulus time-courses, such as sinusoids. Another limitation was that there were only three pulses, which, in terms of experimental design, maps to a paradigm in which conditions only have three trials. In practice, the number of trials per condition will be greater. Therefore, to improve the similarity between results with synthetic data and real-world experiments, further work should be performed with a higher number of pulses. Another limitation was that the sources had non-overlapping time-courses and the spatial maps, apart from the first source in Chapter 2, all comprised single, unique regions. From a biological perspective, it is implausible that functional networks in the brain have activations that are temporally non-overlapping and spatially constrained to single regions. While this was a good starting point to evaluate (p)ICA in 'best-case' conditions, for further investigation, it should be evaluated how (p)ICA performance varies as a function of the temporal and spatial overlap of the sources. Another limitation was that the source-to-region mappings used in Chapters 2 and 4 were fixed across simulations for each experiment. As described in Section 5.1.2, the investigations performed to inspect the effects of the source-to-region mappings were not conclusive and, therefore, for further investigation the experiment described in Section A.2 should be repeated for each region of the first source with multiple source-to-region mappings for the other 'non-target' sources.

In this project, only additive white Gaussian noise, with no temporal autocorrelation, was used to simulate unexplained variance in the LFP signals at the cortical vertex level. While this is commonly used (Deco et al., 2009; Frank et al., 1999; Groves et al., 2011), more sophisticated noise models with more biologically plausible spectral properties (e.g. Rician noise used in Erhardt et al. (2011)) should be used in further investigations. Another

limitation was that all cortical vertices had the same level of noise, and further work could be done to investigate more heterogeneous spatial distributions of noise. Sensor or measurement noise was not added to the EEG or fMRI signals in order to keep parameter space small. For more realistic evaluations of ICA, this should be added in further investigations.

The evaluation of (p)ICA was only performed on the simulated signals of a single subject. For the anatomical features of this subject, such as the connectome and the cortical surface model, the TVB demonstration dataset (Sanz-Leon et al., 2015a) was used. TVB can be used to simulate the data of multiple subjects by using datasets that have different anatomical properties. Originally, we had planned to compare the performances of single-subject (p)ICA with group (p)ICA, but this was not possible as the evaluation of single-modality ICA took much longer than anticipated. These will be useful investigations to perform in the future.

5.4.3 ICA parameters

The number of ICs was always kept to be equal to the number of sources, and the reason for this was to evaluate (p)ICA in 'best-case' conditions (when in the absence of noise). With real data, the true number of ICs is not known and, therefore, model order selection, as described in Section 2.4, is a non-trivial practical issue that needs to be addressed when applying ICA. Therefore, future work that investigates (p)ICA performance as a function of the number of ICs should be informative. It is also possible that ICA performance may improve in the presence of noise when there are more ICs than stimuli. While in earlier investigations, which are not reported in this thesis, no significant differences were seen in ICA performance in the presence of noise with 50 or 51 ICs when there were 50 sources, it will be useful to investigate this more rigourously and systematically.

The only ICA algorithm used in this project was FastICA (Hyvärinen and Oja, 1997), and the main reasons for this were, as described in Section 2.2.3.1, that this is one of the most popular and computationally efficient algorithms. While several comparisons of ICA algorithms have previously been performed on various metrics (e.g. Delorme et al., 2012; Giannakopoulos et al., 1999), to our knowledge, their performances as a function of the number of neural sources and source network size have not been compared. In future investigations, other ICA algorithms could be compared using the framework presented here. Also, FastICA was only evaluated using the tanh nonlinearity and symmetric decomposition (see Section 2.2.3.1), and further investigation could be done at comparing its performance with the other parameters.

To ensure that the inferences on ICA performance were robust across decompostions, in all experiments, ICA was performed 30 times and the performance metric for each simulation was calculated as the mean performance across the decompositions. ICASSO (Himberg and

Hyvarinen, 2003) provides an alternative method for assessing the reliability of the ICs, and is incorporated in some fMRI ICA toolboxes, such as the GIFT toolbox (Rachakonda et al., 2007). For future investigations, ICASSO could also be used and compared with the existing method of the averaging the performance metric across multiple decompositions.

5.4.4 Connectivity measures

Originally, as mentioned in Section 5.1.3, the plan had been to evaluate how well pICA performed in matching the ICs across modalities using multiple connectivity measures. However, due to limits on time, we were unable to perform this. The only measures that were used for pICA were correlation and mutual information. The investigation with mutual information was very limited, as a single toolbox was used, and there are some concerns about the accuracy of the calculations (see Section 4.4.4). The results with mutual information were also not very promising, suggesting that the use of nonlinear measures for pICA might not be useful. From an exploratory point of view, however, evaluating pICA performance with other connectivity measures does make sense, as it is difficult to rule out the possibility that some measures may provide superior performance to correlation. Therefore, for further investigation, pICA could be evaluated with other connectivity measures (see Appendix E).

5.5 Novel contributions

This project makes several novel contributions to the existing body of literature on ICA and its applications for combining EEG and fMRI signals. The primary contribution is the evaluation of 'trial-by-trial' EEG-fMRI pICA (see Section 1.5.1.2), and related to this are devising a basic, 'stripped-down' version of pICA and evaluating its performance as a function of the number of neural sources and level of source noise. Matching ICs across modalities using their spatial features is also a novel idea that is implemented and evaluated, along with the matching of ICs using a nonlinear measure of association (i.e. mutual information). The evaluation of single-modality EEG and fMRI ICA as a function of the number of neural sources and their size are novel contributions as well, as is implicitly testing if EEG ICA and fMRI ICA can provide information about the same sources. One of the key contributions of this project also is that it demonstrates how to generate synthetic data using TVB for the evaluation of unimodal and multimodal neuroimaging methods and how to perform the evaluation. The code and data that were used for this are available on request to the author.

5.6 Conclusions

Human neuroscience has seen a large increase in the number of methods available to analyse neural signals. While it is exciting to apply various signal processing, statistical, and machine learning techniques in novel ways to these signals, it is crucial to not forget that the purpose of the application is to extract accurate biological information. Therefore, method development should be accompanied with the development of evaluation frameworks that can be used to systematically guide the process.

In this project, (p)ICA (EEG ICA, fMRI ICA, and EEG-fMRI parallel ICA) were evaluated in terms of their accuracy at providing information on functional networks in the brain, and the evaluation framework developed for this was also used to test the utility of extending parallel ICA (pICA) by matching ICs across modalities using their spatial features and mutual information. Overall, (p)ICA performance was observed to increase with the number of orthogonal sources in the absence of noise, and to decrease with the level of noise dispersion when the number of sources was fixed. In the absence of noise, EEG and fMRI ICA performances did not vary largely with source network size, but in the presence of noise, they varied without clear trends. The reasons underlying these relationships between (p)ICA performance and the number of sources and single-modality ICA performances and the size of the sources are not currently understood and should be further investigated.

One of the take-home messages from this project is that (p)ICA do not always perform well in ideal conditions, i.e. when the sources are orthogonal and there is no noise. Therefore, while the application of these methods is generally useful, their results should be interpreted cautiously and validated where possible. Two other take-home messages are that pICA performs better when spatial features of the ICs are used to match them across modalities, in addition to their temporal features, and that mutual information does not appear to be useful for matching the ICs and performs significantly worse than correlation.

It must be acknowleged that the results reported in this project and the conclusions drawn are based on initial investigations that apply a novel evaluation framework that uses The Virtual Brain (Sanz-Leon et al., 2013) for generating synthetic data. Further work should be performed with this framework that uses more diverse instrinsic and extrinstic simulation parameters to infer on, and to improve, the generalisability of the observations.

References

Abou-Elseoud, A., Starck, T., Remes, J., Nikkinen, J., Tervonen, O., and Kiviniemi, V. (2010). The effect of model order selection in group PICA. *Human Brain Mapping*, 31(8):1207–1216. 00290.

Abreu, R., Leal, A., and Figueiredo, P. (2018). EEG-Informed fMRI: A Review of Data Analysis Methods. *Frontiers in Human Neuroscience*, 12.

Amari, S.-i., Cichocki, A., and Yang, H. H. (1996). A New Learning Algorithm for Blind Signal Separation. In Touretzky, D. S., Mozer, M. C., and Hasselmo, M. E., editors, *Advances in Neural Information Processing Systems 8*, pages 757–763. MIT Press. 02828.

Ashby, F. G. (2011). *Statistical analysis of fMRI data*. MIT press.

Assecondi, S., Ostwald, D., and Bagshaw, A. P. (2015). Reliability of Information-Based Integration of EEG and fMRI Data: A Simulation Study. *Neural Computation*, 27(2):281–305.

Bagshaw, A. P. and Warbrick, T. (2007). Single trial variability of EEG and fMRI responses to visual stimuli. *NeuroImage*, 38(2):280–292.

Bakker, R., Wachtler, T., and Diesmann, M. (2012). CoCoMac 2.0 and the future of tract-tracing databases. *Frontiers in Neuroinformatics*, 6. 00060.

Beckmann, C. F. (2012). Modelling with independent components. *NeuroImage*, 62(2):891–901.

Beckmann, C. F. and Smith, S. M. (2005). Tensorial extensions of independent component analysis for multisubject FMRI analysis. *NeuroImage*, 25(1):294–311. 00755.

Bell, A. J. and Sejnowski, T. J. (1995). An Information-Maximization Approach to Blind Separation and Blind Deconvolution. *Neural Computation*, 7(6):1129–1159.

Bridwell, D. A., Rachakonda, S., Silva, R. F., Pearlson, G. D., and Calhoun, V. D. (2016). Spatiospectral Decomposition of Multi-subject EEG: Evaluating Blind Source Separation Algorithms on Real and Realistic Simulated Data. *Brain Topography*.

Bridwell, D. A., Wu, L., Eichele, T., and Calhoun, V. D. (2013). The spatiospectral characterization of brain networks: Fusing concurrent EEG spectra and fMRI maps. *NeuroImage*, 69:101–111.

Bénar, C.-G., Schön, D., Grimault, S., Nazarian, B., Burle, B., Roth, M., Badier, J.-M., Marquis, P., Liegeois-Chauvel, C., and Anton, J.-L. (2007). Single-trial analysis of oddball event-related potentials in simultaneous EEG-fMRI. *Human Brain Mapping*, 28(7):602–613.

Calhoun, V., Adali, T., Giuliani, N., Pekar, J., Kiehl, K., and Pearlson, G. (2006a). Method for multimodal analysis of independent source differences in schizophrenia: Combining gray matter structural and auditory oddball functional data. *Human Brain Mapping*, 27(1):47–62.

Calhoun, V., Adali, T., McGinty, V., Pekar, J., Watson, T., and Pearlson, G. (2001a). fMRI Activation in a Visual-Perception Task: Network of Areas Detected Using the General Linear Model and Independent Components Analysis. *NeuroImage*, 14(5):1080–1088.

Calhoun, V., Adali, T., Pearlson, G., and Kiehl, K. (2006b). Neuronal chronometry of target detection: Fusion of hemodynamic and event-related potential data. *NeuroImage*, 30(2):544–553.

Calhoun, V., Adali, T., Pearlson, G., and Pekar, J. (2001b). A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping*, 14(3):140–151.

Calhoun, V. D., Adali, T., Hansen, L. K., Larsen, J., and Pekar, J. J. (2003). ICA of Functional MRI Data: An Overview. In *in Proceedings of the International Workshop on Independent Component Analysis and Blind Signal Separation*, pages 281–288.

Calhoun, V. D., Adali, T., Pearlson, G. D., and Pekar, J. J. (2001c). Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Human brain mapping*, 13(1):43–53.

Calhoun, V. D., Kiehl, K. A., and Pearlson, G. D. (2008). Modulation of temporally coherent brain networks estimated using ICA at rest and during cognitive tasks. *Human Brain Mapping*, 29(7):828–838.

Calhoun, V. D. and Lacy, N. d. (2017). Ten Key Observations on the Analysis of Resting-state Functional MR Imaging Data Using Independent Component Analysis. *Neuroimaging Clinics*, 27(4):561–579. 00000.

Calhoun, V. D., Liu, J., and Adalı, T. (2009). A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *NeuroImage*, 45(1):S163–S172.

Capilla, A., Melcón, M., Kessel, D., Calderón, R., Pazo-Álvarez, P., and Carretié, L. (2016). Retinotopic mapping of visual event-related potentials. *Biological Psychology*, 118:114–125.

Correa, N., Adalı, T., and Calhoun, V. D. (2007). Performance of blind source separation algorithms for fMRI analysis using a group ICA method. *Magnetic Resonance Imaging*, 25(5):684–694.

Cover, T. M. and Thomas, J. A. (2006). *Elements of information theory*. Wiley-Interscience, Hoboken, N.J, 2nd ed edition.

Daunizeau, J., David, O., and Stephan, K. (2011). Dynamic causal modelling: A critical review of the biophysical and statistical foundations. *NeuroImage*, 58(2):312–322.

De Luca, M., Beckmann, C. F., De Stefano, N., Matthews, P. M., and Smith, S. M. (2006). fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *NeuroImage*, 29(4):1359–1367.

Deco, G., Jirsa, V., McIntosh, A. R., Sporns, O., and Kötter, R. (2009). Key role of coupling, delay, and noise in resting brain fluctuations. *Proceedings of the National Academy of Sciences*, 106(25):10302–10307.

Deco, G., Jirsa, V. K., Robinson, P. A., Breakspear, M., and Friston, K. (2008). The Dynamic Brain: From Spiking Neurons to Neural Masses and Cortical Fields. *PLoS Computational Biology*, 4(8):e1000092.

Delorme, A. and Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134(1):9–21.

Delorme, A., Palmer, J., Onton, J., Oostenveld, R., Makeig, S., and others (2012). Independent EEG sources are dipolar. *PloS one*, 7(2):e30135.

Edwards, B. G., Calhoun, V. D., and Kiehl, K. A. (2012). Joint ICA of ERP and fMRI during error-monitoring. *NeuroImage*, 59(2):1896–1903.

Eichele, T. and Calhoun, V. D. (2010). Parallel EEG-fMRI ICA Decomposition. *Simultaneous EEG and fMRI: Recording, Analysis, and Application*, page 175.

Eichele, T., Calhoun, V. D., Moosmann, M., Specht, K., Jongsma, M. L., Quiroga, R. Q., Nordby, H., and Hugdahl, K. (2008). Unmixing concurrent EEG-fMRI with parallel independent component analysis. *International Journal of Psychophysiology*, 67(3):222–234.

Erhardt, E. B., Allen, E. A., Wei, Y., Eichele, T., and Calhoun, V. D. (2012). SimTB, a simulation toolbox for fMRI data under a model of spatiotemporal separability. *NeuroImage*, 59(4):4160–4167. 00146.

Erhardt, E. B., Rachakonda, S., Bedrick, E. J., Allen, E. A., Adali, T., and Calhoun, V. D. (2011). Comparison of multi-subject ICA methods for analysis of fMRI data. *Human Brain Mapping*, 32(12):2075–2095.

Faisal, A. A., Selen, L. P. J., and Wolpert, D. M. (2008). Noise in the nervous system. *Nature reviews. Neuroscience*, 9(4):292–303.

Fisch, B. J. and Spehlmann, R. (1999). *Fisch and Spehlmann's EEG primer: basic principles of digital and analog EEG*. Elsevier Health Sciences.

Frank, T. D., Daffertshofer, A., Beek, P. J., and Haken, H. (1999). Impacts of noise on a field theoretical model of the human brain. *Physica D: Nonlinear Phenomena*, 127(3):233–249.

Friston, K., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4):1273–1302.

Friston, K., Mechelli, A., Turner, R., and Price, C. (2000). Nonlinear Responses in fMRI: The Balloon Model, Volterra Kernels, and Other Hemodynamics. *NeuroImage*, 12(4):466–477.

Friston, K., Moran, R., and Seth, A. K. (2013). Analysing connectivity with Granger causality and dynamic causal modelling. *Current opinion in neurobiology*, 23(2):172–178.

Friston, K. J. (2011). Functional and Effective Connectivity: A Review. *Brain Connectivity*, 1(1):13–36.

Ghahremani, D., Makeig, S., Jung, T., Bell, A. J., and Sejnowski, T. J. (1996). Independent Component Analysis of Simulated EEG Using a Three-Shell Spherical Head Model. Technical report, DTIC Document.

Giannakopoulos, X., Karhunen, J., and Oja, E. (1999). An experimental comparison of neural algorithms for independent component analysis and blind separation. *International journal of neural systems*, 9(02):99–114.

Gjedde, A. (2001). Brain energy metabolism and the physiological basis of the haemodynamic response. *Functional MRI : an introduction to methods*, pages 38–67.

Goense, J., Logothetis, N. K., and Debener, U. S. (2010). Physiological basis of the BOLD signal. *Simultaneous EEG and fMRI: Recording, analysis, and application*, pages 21–46.

Gonçalves, S., de Munck, J., Pouwels, P., Schoonhoven, R., Kuijer, J., Maurits, N., Hoogduin, J., Van Someren, E., Heethaar, R., and Lopes da Silva, F. (2006). Correlating the alpha rhythm to BOLD using simultaneous EEG/fMRI: Inter-subject variability. *NeuroImage*, 30(1):203–213.

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, 37(3):424.

Greicius, M. D., Flores, B. H., Menon, V., Glover, G. H., Solvason, H. B., Kenna, H., Reiss, A. L., and Schatzberg, A. F. (2007). Resting-State Functional Connectivity in Major Depression: Abnormally Increased Contributions from Subgenual Cingulate Cortex and Thalamus. *Biological Psychiatry*, 62(5):429–437.

Groves, A. R., Beckmann, C. F., Smith, S. M., and Woolrich, M. W. (2011). Linked independent component analysis for multimodal data fusion. *NeuroImage*, 54(3):2198–2217.

Gävert, H., Hurri, J., Särelä, J., and Hyvärinen, A. (2005). FastICA for Matlab. https://research.ics.aalto.fi/ica/fastica/code/FastICA_2.5.zip, last accessed on 01/05/2019.

Himberg, J. and Hyvarinen, A. (2003). Icasso: software for investigating the reliability of ICA estimates by clustering and visualization. In *2003 IEEE XIII Workshop on Neural Networks for Signal Processing (IEEE Cat. No.03TH8718)*, pages 259–268. 00196 ISSN: 1089-3555.

Hui, M., Li, J., Wen, X., Yao, L., and Long, Z. (2011). An Empirical Comparison of Information-Theoretic Criteria in Estimating the Number of Independent Components of fMRI Data. *PLOS ONE*, 6(12):e29274. 00016.

Huster, R. J., Debener, S., Eichele, T., and Herrmann, C. S. (2012). Methods for Simultaneous EEG-fMRI: An Introductory Review. *Journal of Neuroscience*, 32(18):6053–6060.

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., Della Penna, S., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J., Handwerker, D. A., Keilholz, S., Kiviniemi, V., Leopold, D. A., de Pasquale, F., Sporns, O., Walter, M., and Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage*, 80:360–378.

Hyvärinen, A., Karhunen, J., and Oja, E. (2004). *Independent component analysis*, volume 46. John Wiley & Sons.

Hyvärinen, A. and Oja, E. (1997). A Fast Fixed-Point Algorithm for Independent Component Analysis. *Neural Computation*, 9(7):1483–1492. 04273 Publisher: MIT Press.

Hyvärinen, A. and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural networks*, 13(4):411–430.

Ives, J. R., Warach, S., Schmitt, F., Edelman, R. R., and Schomer, D. L. (1993). Monitoring the patient's EEG during echo planar MRI. *Electroencephalography and Clinical Neurophysiology*, 87(6):417–420.

Jann, K., Dierks, T., Boesch, C., Kottlow, M., Strik, W., and Koenig, T. (2009). BOLD correlates of EEG alpha phase-locking and the fMRI default mode network. *NeuroImage*, 45(3):903–916.

Jonmohamadi, Y., Poudel, G., Innes, C., and Jones, R. (2014). Source-space ICA for EEG source separation, localization, and time-course reconstruction. *NeuroImage*, 101(Supplement C):720–737. 00016.

Kantz, H. and Schreiber, T. (2004). *Nonlinear time series analysis*, volume 7. Cambridge university press.

Kozachenko, L. F. and Leonenko, N. N. (1987). Sample estimate of the entropy of a random vector. *Problemy Peredachi Informatsii*, 23(2):9–16.

Kraskov, A., Stögbauer, H., and Grassberger, P. (2004). Estimating mutual information. *Physical review E*, 69(6):066138.

Laufs, H. (2012). A personalized history of EEG–fMRI integration. *NeuroImage*, 62(2):1056–1067.

Laufs, H., Kleinschmidt, A., Beyerle, A., Eger, E., Salek-Haddadi, A., Preibisch, C., and Krakow, K. (2003). EEG-correlated fMRI of human alpha activity. *NeuroImage*, 19(4):1463–1476.

Lee, T.-W., Girolami, M., Bell, A. J., and Sejnowski, T. J. (2000). A unifying information-theoretic framework for independent component analysis. *Computers & Mathematics with Applications*, 39(11):1–21.

Lei, X., Qiu, C., Xu, P., and Yao, D. (2010). A parallel framework for simultaneous EEG/fMRI analysis: Methodology and simulation. *NeuroImage*, 52(3):1123–1134.

Lei, X., Valdes-Sosa, P. A., and Yao, D. (2012). EEG/fMRI fusion based on independent component analysis: Integration of data-driven and model-driven methods. *Journal of Integrative Neuroscience*, 11(03):313–337.

Li, Y.-O., Adalı, T., and Calhoun, V. D. (2007). Estimating the number of independent components for functional magnetic resonance imaging data. *Human Brain Mapping*, 28(11):1251–1266.

Lindner, M., Vicente, R., Priesemann, V., and Wibral, M. (2011). TRENTOOL: A Matlab open source toolbox to analyse information flow in time series data with transfer entropy. *BMC neuroscience*, 12(1):119.

Liu, J. and Calhoun, V. (2007). Parallel independent component analysis for multimodal analysis: application to fMRI and EEG data. In *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1028–1031. IEEE.

Liu, J., Pearlson, G., Windemuth, A., Ruano, G., Perrone-Bizzozero, N. I., and Calhoun, V. (2009). Combining fMRI and SNP data to investigate connections between brain function and genetics using parallel ICA. *Human Brain Mapping*, 30(1):241–255.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197):869–878.

Mainen, Z. and Sejnowski, T. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506.

Makeig, S. (2002). Dynamic Brain Sources of Visual Evoked Responses. *Science*, 295(5555):686–690.

Makeig, S., Bell, A. J., Jung, T.-P., Sejnowski, T. J., and others (1996). Independent component analysis of electroencephalographic data. *Advances in neural information processing systems*, pages 145–151.

Makeig, S.,, Jung, T. P., Bell, A. J., Ghahremani, D.,, and Sejnowski, T. J. (1997). Blind separation of auditory event-related brain responses into independent components. Proceedings of the National. *Proceedings of the National Academy of Sciences*, 94(20):10979–10984.

Materassi, D. (2014). Norbert Wiener's legacy in the study and inference of causation. In *Norbert Wiener in the 21st Century (21CW), 2014 IEEE Conference on*, pages 1–6. IEEE.

Moosmann, M., Eichele, T., Nordby, H., Hugdahl, K., and Calhoun, V. D. (2008). Joint independent component analysis for simultaneous EEG–fMRI: Principle and simulation. *International Journal of Psychophysiology*, 67(3):212–221.

Novitskiy, N., Ramautar, J., Vanderperren, K., De Vos, M., Mennes, M., Mijovic, B., Vanrumste, B., Stiers, P., Van den Bergh, B., Lagae, L., Sunaert, S., Van Huffel, S., and Wagemans, J. (2011). The BOLD correlates of the visual P1 and N1 in single-trial analysis of simultaneous EEG-fMRI recordings during a spatial detection task. *NeuroImage*, 54(2):824–835.

Nunez, P. L. and Silberstein, R. B. (2000). On the relationship of synaptic activity to macroscopic measurements: does co-registration of EEG with fMRI make sense? *Brain topography*, 13(2):79–96.

Ogawa, S., Lee, T. M., Kay, A. R., and Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87(24):9868–9872.

Oostenveld, R., Fries, P., Maris, E., and Schoffelen, J.-M. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. 04489 ISSN: 1687-5265 Library Catalog: www.hindawi.com Pages: e156869 Publisher: Hindawi Volume: 2011.

Pascual-Marqui, R. D. (1999). Review of methods for solving the EEG inverse problem. *International journal of bioelectromagnetism*, 1(1):75–86.

Peng, H., Long, F., and Ding, C. (2005). Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238. 07486 Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.

Rachakonda, S., Egolf, E., Correa, N., and Calhoun, V. (2007). Group ICA of fMRI toolbox (GIFT) manual. *Dostupné z: http://www. nitrc. org/docman/view. php/55/295/v1. 3d_GIFTManual. pdf [cit. 2011-11-5]*.

Rachakonda, S., Liu, J., and Calhoun, V. (2008). Fusion ICA toolbox (FIT) manual. *Albuquerque, NM: The MIND Research Network, University of New Mexico*.

Risk, B. B., Matteson, D. S., Ruppert, D., Eloyan, A., and Caffo, B. S. (2014). An evaluation of independent component analyses with an application to resting-state fMRI. *Biometrics*, 70(1):224–236.

Rosa, M. J., Daunizeau, J., and Friston, K. J. (2010). EEG-fMRI integration: a critical review of biophysical modelling and data analysis approaches. *Journal of Integrative Neuroscience*, 09(04):453–476.

Sanz-Leon, P., Knock, S. A., Spiegler, A., and Jirsa, V. K. (2015a). A Description of a Complete Dataset. http://docs.thevirtualbrain.org/manuals/UserGuide/Complete_Dataset_Description.html, last accessed on 01/05/2019.

Sanz-Leon, P., Knock, S. A., Spiegler, A., and Jirsa, V. K. (2015b). Mathematical framework for large-scale brain network modeling in The Virtual Brain. *NeuroImage*, 111:385–430. 00055.

Sanz-Leon, P., Knock, S. A., Woodman, M. M., Domide, L., Mersmann, J., McIntosh, A. R., and Jirsa, V. (2013). The Virtual Brain: a simulator of primate brain network dynamics. *Frontiers in Neuroinformatics*, 7.

Scheeringa, R., Petersson, K. M., Oostenveld, R., Norris, D. G., Hagoort, P., and Bastiaansen, M. C. (2009). Trial-by-trial coupling between EEG and BOLD identifies networks related to alpha and theta EEG power increases during working memory maintenance. *NeuroImage*, 44(3):1224–1238.

Schreiber, T. (2000). Measuring information transfer. *Physical review letters*, 85(2):461.

Schürmann, T. (2004). Bias analysis in entropy estimation. *Journal of Physics A: Mathematical and General*, 37(27):L295–L301.

Seth, A. K., Barrett, A. B., and Barnett, L. (2015). Granger Causality Analysis in Neuroscience and Neuroimaging. *Journal of Neuroscience*, 35(8):3293–3297.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3):379–423.

Strobel, A., Debener, S., Sorger, B., Peters, J. C., Kranczioch, C., Hoechstetter, K., Engel, A. K., Brocke, B., and Goebel, R. (2008). Novelty and target processing during an auditory novelty oddball: A simultaneous event-related potential and functional magnetic resonance imaging study. *NeuroImage*, 40(2):869–883.

Valdes-Sosa, P. A., Sanchez-Bornot, J. M., Sotero, R. C., Iturria-Medina, Y., Aleman-Gomez, Y., Bosch-Bayard, J., Carbonell, F., and Ozaki, T. (2009). Model driven EEG/fMRI fusion of brain oscillations. *Human Brain Mapping*, 30(9):2701–2721.

Vanni, S., Warnking, J., Dojat, M., Delon-Martin, C., Bullier, J., and Segebarth, C. (2004). Sequence of pattern onset responses in the human visual areas: an fMRI constrained VEP source analysis. *NeuroImage*, 21(3):801–817. 00000.

Vergara, V. M., Ulloa, A., Calhoun, V. D., Boutte, D., Chen, J., and Liu, J. (2014). A three-way parallel ICA approach to analyze links among genetics, brain structure and brain function. *NeuroImage*, 98:386–394.

Vicente, R., Wibral, M., Lindner, M., and Pipa, G. (2011). Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of Computational Neuroscience*, 30(1):45–67.

Vigario, R., Sarela, J., Jousmiki, V., Hamalainen, M., and Oja, E. (2000). Independent component approach to the analysis of EEG and MEG recordings. *IEEE Transactions on Biomedical Engineering*, 47(5):589–593. 00687.

Vigário, R. and Oja, E. (2000). Independence: a new criterion for the analysis of the electromagnetic fields in the global brain? *Neural networks*, 13(8):891–907.

Wei, H., Jafarian, A., Zeidman, P., Litvak, V., Razi, A., Hu, D., and Friston, K. J. (2020). Bayesian fusion and multimodal DCM for EEG and fMRI. *NeuroImage*, 211:116595. 00002.

Wibral, M., Turi, G., Linden, D. E. J., Kaiser, J., and Bledowski, C. (2008). Decomposition of working memory-related scalp ERPs: Crossvalidation of fMRI-constrained source analysis and ICA. *International Journal of Psychophysiology*, 67(3):200–211. 00022.

Wiener, N. (1956). The theory of prediction. *Modern mathematics for engineers*, 1:125–139.

Wilson, H. R. and Cowan, J. D. (1972). Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons. *Biophysical Journal*, 12(1):1–24.

Wong, K.-F. and Wang, X.-J. (2006). A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *Journal of Neuroscience*, 26(4):1314–1328.

Wu, L., Eichele, T., and Calhoun, V. D. (2010). Reactivity of hemodynamic responses and functional connectivity to different states of alpha synchrony: A concurrent EEG-fMRI study. *NeuroImage*, 52(4):1252–1260.

Yoshioka, T., Toyama, K., Kawato, M., Yamashita, O., Nishina, S., Yamagishi, N., and Sato, M.-a. (2008). Evaluation of hierarchical Bayesian method through retinotopic brain activities reconstruction from fMRI and MEG signals. *NeuroImage*, 42(4):1397–1413.

Appendix A

Chapter 2: Further Investigations with Source-to-Region Mappings

A.1 Introduction

In this appendix, two further investigations with source-to-region mappings that were performed after the completion of Chapter 2 are reported. The experiment numbers in this appendix correspond to those in Chapter 2, e.g. Experiment 1 refers to Experiment 1 described in Section 2.2.4.1.

In the first investigation (Section A.2), to check if the region the first source comprised affected ICA performance, an experiment was run in which the source comprised each region individually across simulations. It was concluded from this experiment that EEG ICA performance varied not only depending on the region the source comprised but also depending on the regions the other sources comprised as well. In the second investigation (Section A.3), to check if the results of Experiment 1 contained the dip in fMRI ICA performance when the source-to-region mappings were varied for all sources, it was repeated an additional nine times, and it was observed that in most cases the dip was visible and the mean trend across the replications was similar to what had been observed originally in the experiment.

In Section A.4 it is discussed that these investigations raise further questions as they are somewhat contradictory. The first investigation shows that the source-to-region mappings did affect EEG ICA performance, whereas the second investigation shows that the source-to-region mappings did not largely affect ICA performance in either modality. It is suggested that the anatomical configuration of all sources, rather than just the first one, may affect the EEG ICA performance of the first source, and further work is needed to investigate this.

A.2 Does ICA performance depend on which region the first source comprises?

A.2.1 Introduction

In Experiment 1, EEG ICA performance is 0.843 when there are 50 sources and the size of the source is 1 region (Figure 2.8). However, in Experiment 2, when there are 50 sources and the size of the source is 1 region, EEG ICA performance is 0.304 (Figure 2.9). In these two experiments, the region that the first source comprises is different, suggesting that ICA performance varies depending on the anatomical location of the source. In this study, it had been assumed that this was not the case, as ICA separates sources based on their statistical independence and is blind to underlying forward models. To test this assumption, 76 simulations were run, one for each region, to see if ICA performance varied with the location of the first source.

A.2.2 Methods

76 simulations were run. In each simulation, there were 50 sources. The source-to-region mapping of the last 49 sources remained constant across the simulations, and, for these sources, the mapping used in Experiments 2, 5 and 6 was used (Figure 2.5). The source-to-region mapping of the first source systematically varied across the simulations - in the first simulation it only comprised the first region, in the second simulation it only comprised the second region, and so on. ICA performance for both modalities was calculated using the method described in the main text of the manuscript (see Section 2.2.3.2).

A.2.3 Results

EEG and fMRI ICA performances as functions of the region the first source comprised as it varied from region 1 to region 76 are shown in Figure A.1.

A.2.4 Discussion

A.2.4.1 Effect of anatomical location of the first source

The performance of fMRI ICA is high (> 0.8) irrespective of the region the first source comprises. The performance of EEG ICA, however, varies significantly depending on the region. Table A.1 lists the regions provided in the default TVB parcellation (Sanz-Leon et al., 2015a). The three regions that had the lowest EEG ICA performance were the left

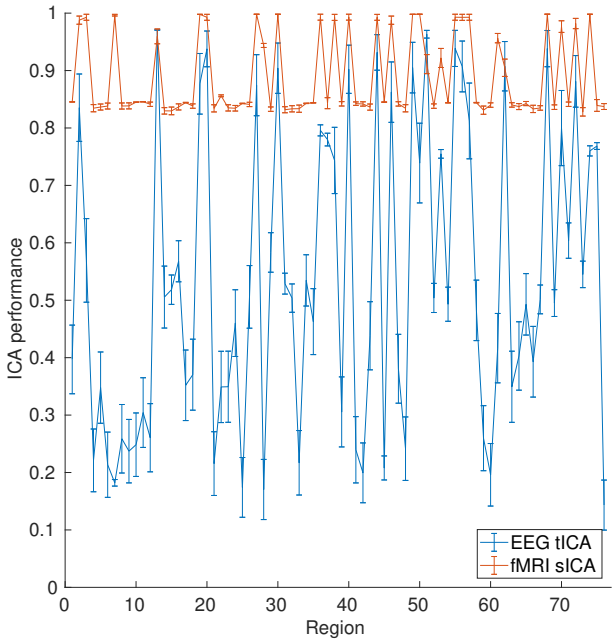

Figure A.1 ICA performance as the region the first source comprised varied from 1 to 76. All sources comprised single, unique regions. Each point represents the mean ICA performance across 30 decompositions, and error bars indicate standard errors of the means.

cingulate cortex (region 76), the right ventrolateral premotor cortex (region 28), and the right parahippocampal cortex (region 25). The three regions that had the highest EEG ICA performance were the right primary motor area (region 13), the left primary motor area (region 51), and the left superior parietal cortex (region 55). To test if regions that are closer to the scalp (and therefore more likely to contribute towards EEG) have better EEG ICA performance, the correlation between the magnitude of the total contribution of each region to the EEG with its EEG ICA performance was calculated. The region-to-EEG channel mapping was constructed by multiplying the region-to-vertex mapping matrix (regions x vertices) with the EEG projection matrix (vertices x channels). The mapping is shown in Figure A.2, and a scatterplot of the relationship is shown in Figure A.3. The strength of the correlation (r) was 0.318 ($p = 0.0051$). Contrary to what had been expected, it can be seen in Figure A.3 that EEG contribution and EEG ICA performance are not strongly related, e.g. there are multiple regions that have high performance (> 0.7) and low contribution ($< 0.5 \times 10^5$). Therefore, it is unclear why EEG ICA performance varies for different regions and this should be further investigated. One reason for this could be the interaction between the region the first source comprises and the source-to-region mappings of the other sources.

Table A.1 TVB regions

Number	Label	Number	Label
1	rA1	39	lA1
2	rA2	40	lA2
3	rAMYG	41	lAMYG
4	rCCA	42	lCCA
5	rCCP	43	lCCP
6	rCCR	44	lCCR
7	rCCS	45	lCCS
8	rFEF	46	lFEF
9	rG	47	lG
10	rHC	48	lHC
11	rIA	49	lIA
12	rIP	50	lIP
13	rM1	51	lM1
14	rPCI	52	lPCI
15	rPCIP	53	lPCIP
16	rPCM	54	lPCM
17	rPCS	55	lPCS
18	rPFCCL	56	lPFCCL
19	rPFCDL	57	lPFCDL
20	rPFCDM	58	lPFCDM
21	rPFCM	59	lPFCM
22	rPFCORB	60	lPFCORB
23	rPFCPOL	61	lPFCPOL
24	rPFCVL	62	lPFCVL
25	rPHC	63	lPHC
26	rPMCDL	64	lPMCDL
27	rPMCM	65	lPMCM
28	rPMCVL	66	lPMCVL
29	rS1	67	lS1
30	rS2	68	lS2
31	rTCC	69	lTCC
32	rTCI	70	lTCI
33	rTCPOL	71	lTCPOL
34	rTCS	72	lTCS
35	rTCV	73	lTCV
36	rV1	74	lV1
37	rV2	75	lV2
38	rCC	76	lCC

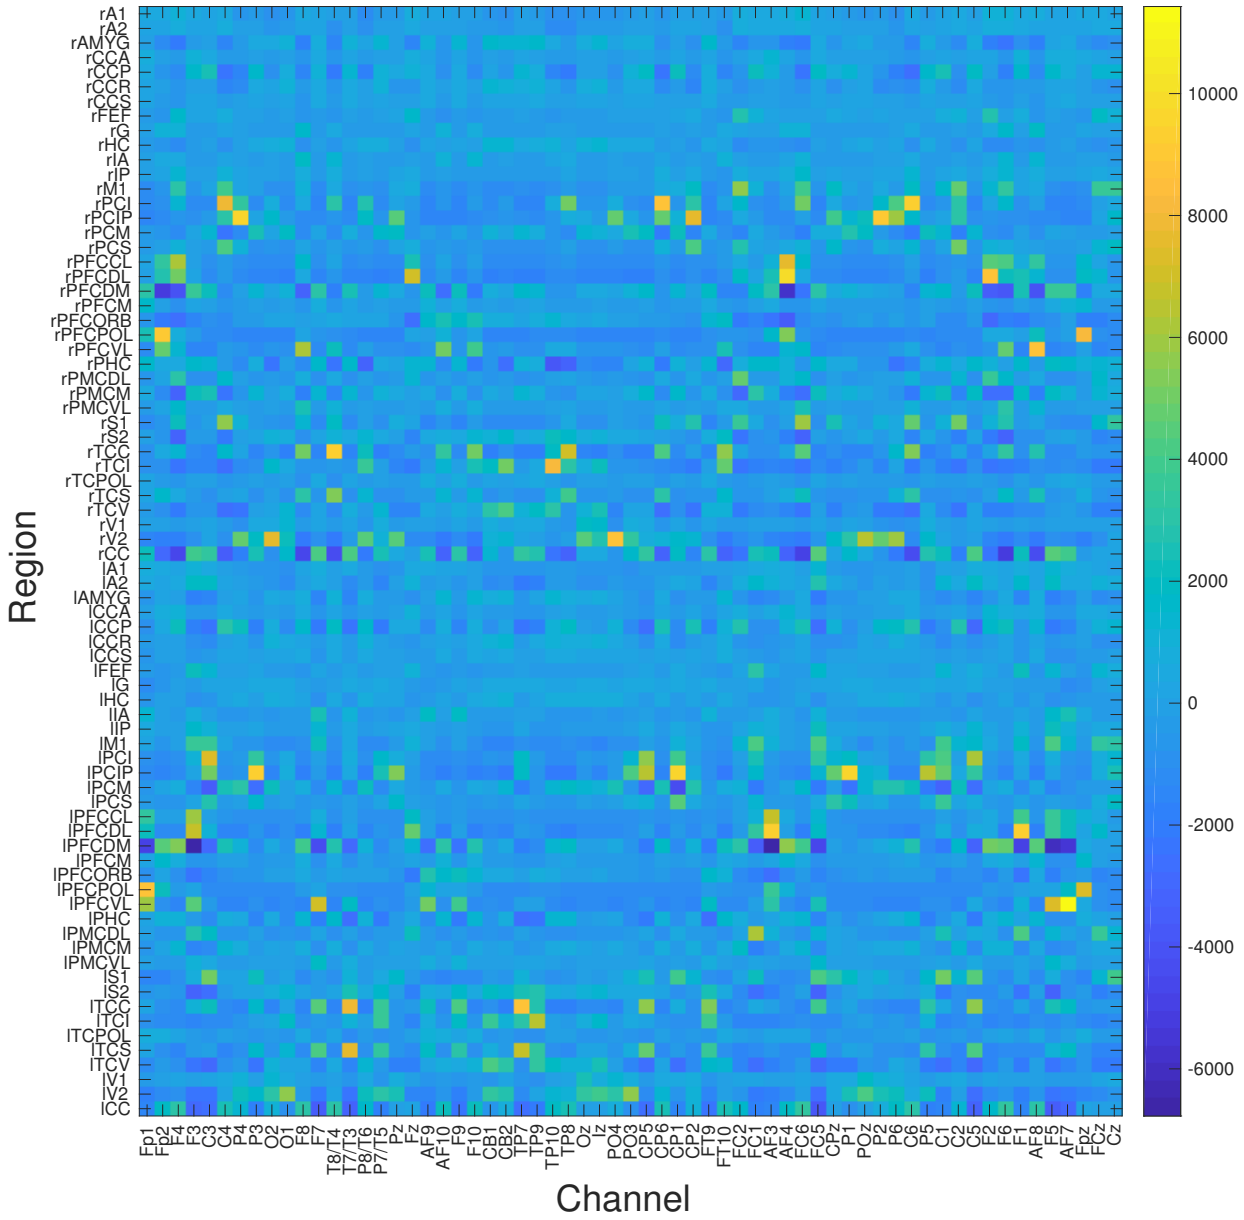

Figure A.2 Region-to-EEG channel mapping

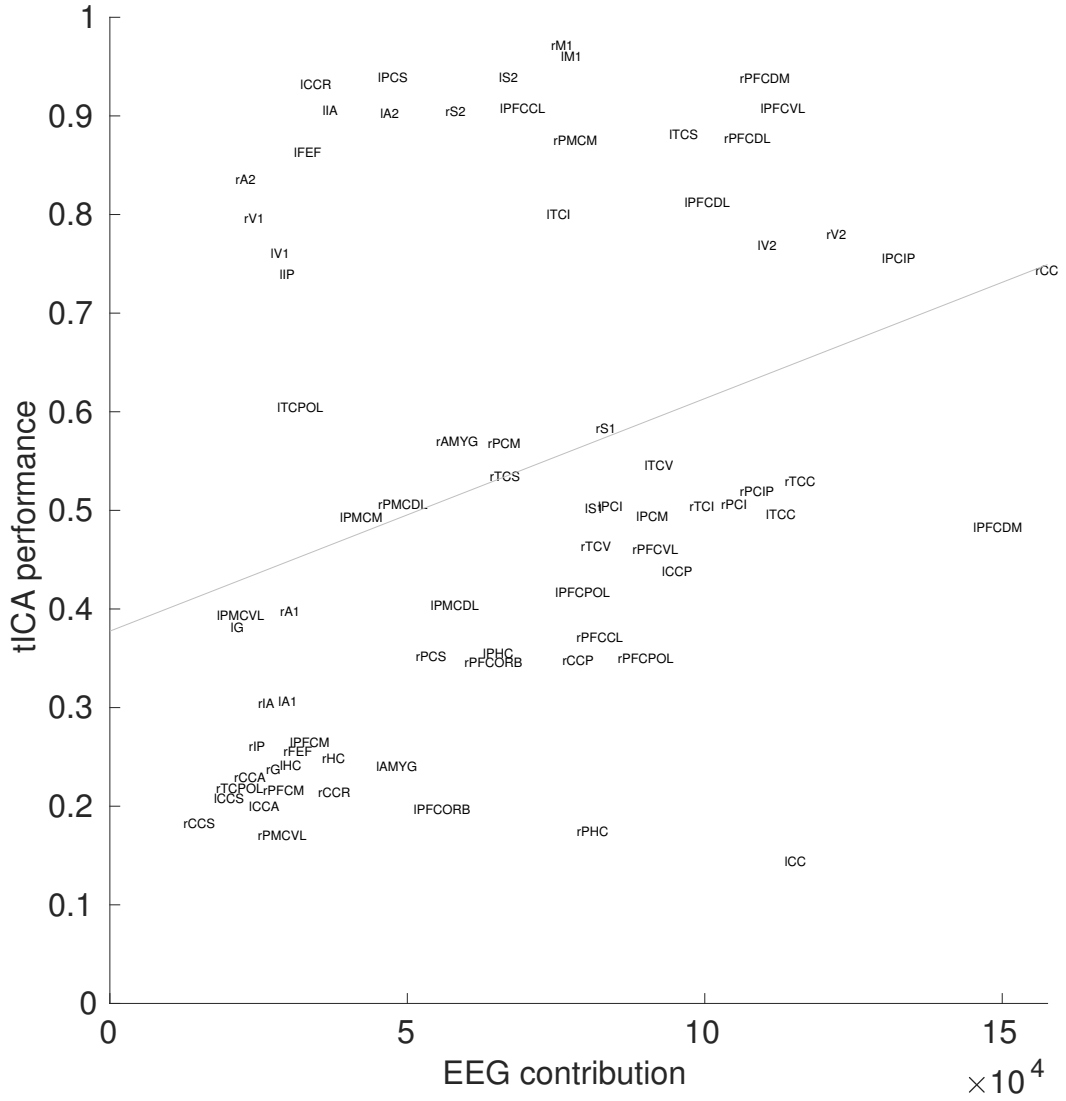

Figure A.3 The correlation between EEG ICA performance and EEG contribution across regions was $r = 0.318$. (The locations of the some of the region labels have been slightly adjusted to prevent the overlap of text.)

A.2.4.2 Consistency with previous experiments

In Experiment 1, the first source comprised region 66 and had EEG ICA performance of 0.843, whereas in Experiment 2, when the size of the first source was 1 region, it comprised region 1 and had EEG ICA performance of 0.304. In this experiment, EEG ICA performance is not completely consistent with these results: while region 1 has EEG ICA performance of 0.397, which is similar to what was observed in Experiment 2, region 66 has EEG ICA performance of 0.393, which is significantly lower to what was observed in Experiment 1. One reason for this could be that the source-to-region mappings of the other 49 sources used were from Experiment 2 and not Experiment 1, suggesting that EEG ICA performance is not only dependent on the region the first source comprises, but rather on the regions all the sources comprise. This question, however, was addressed in Experiment 3, where it was shown that EEG ICA performance did not largely vary across different source-to-region mappings of the non-target source. The inconsistency, however, between the results of Experiment 3 and the observations here might be due to the fact that in Experiment 3 the first source comprised 37 regions, and therefore had a higher signal-to-noise ratio at the sensor-level, whereas here it only comprised a single region.

A.3 Do the results of Experiment 1 vary across source-to-region mappings?

A.3.1 Introduction

The dip in fMRI ICA performance in Experiment 1 (Figure 2.8) is unusual. Here, Experiment 1 was run nine more times with different source-to-region mappings to see if this observation is robust.

A.3.2 Methods

Experiment 1 was run nine more times with different source-to-region mappings, in which each source comprised a single, unique region randomly drawn from a uniform distribution without replacement, and EEG and fMRI ICA performances were calculated for each iteration as described in Section 2.2.3.2. The mean EEG and fMRI ICA performances across the iterations was also calculated.

A.3.3 Results

The individual EEG ICA and fMRI ICA performances are shown in Figure A.4. The mean across iterations is shown in Figure A.5.

A.3.4 Discussion

Across the source-to-region mappings, EEG ICA performance always has a 'mostly' linear increasing trend, whereas fMRI ICA performance is more variable. For example, in Figure A.4a and Figure A.4b, there is an early dip in fMRI ICA performance, whereas in A.4c fMRI ICA performance is always high (> 0.88). In the mean across iterations, EEG ICA performance linearly increases and fMRI ICA performance retains the early dip, somewhat similarly to what was originally reported for Experiment 1 (Figure 2.8) in Chapter 2.

A.4 Conclusions

The results of the experiment performed in the first investigation that looks at how EEG and fMRI ICA performances vary based on the anatomical location of the first source suggest that the assumption that the anatomical locations of the sources would not largely affect ICA performance was incorrect. The results of the second investigation, however, suggest that while there are variations in EEG and fMRI ICA performances with the source-to-region mappings, the overall trend remains largely unchanged as the number of sources increases, with EEG ICA performance linearly increasing and fMRI ICA performance having an early dip. In terms of EEG ICA performance, these results are somewhat contradictory, as the first investigation shows that it was affected by the source-to-region mappings, whereas the second one shows that, in most cases, the trend with the number of sources was largely unaffected. What these results suggest is that the anatomical configuration of all 50 sources may play a role in the ICA performance of the first source, rather than just the location of the first source, and this, perhaps, may be due to underlying local and long-range connectivity. Therefore, to test this, as a next step, it will be useful to repeat the experiment in Section A.2 for each region the first source comprises with multiple source-to-region mappings for the other 49 sources, and then see how much variability there is in the performance for each region. This investigation has not been performed yet as it is computationally very demanding, as each simulation takes approximately 50 hrs to run and analyse.

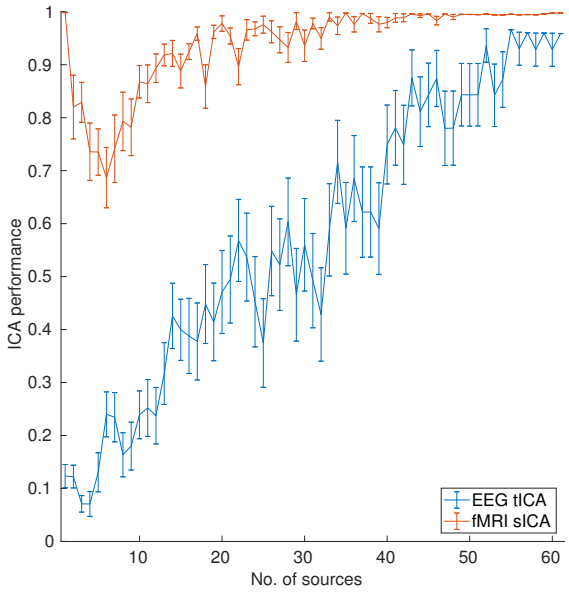

(a) Iteration 1 (reported in main text, Figure 2.8)

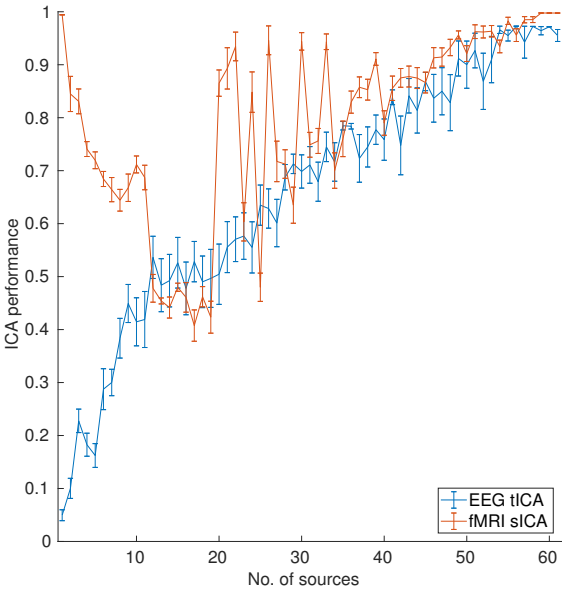

(b) Iteration 2

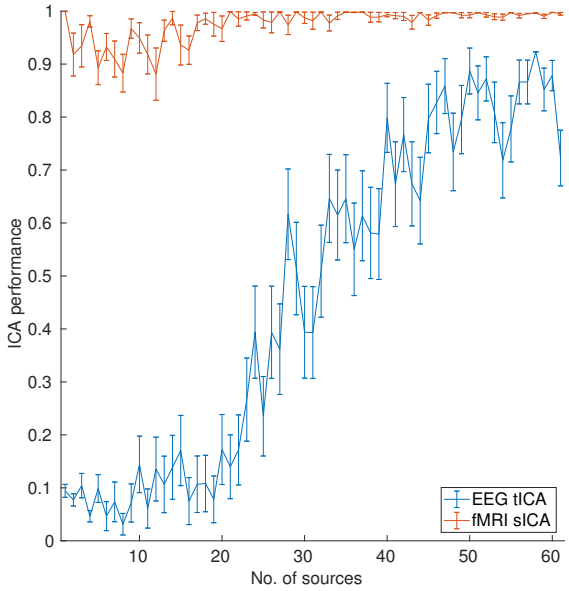

(c) Iteration 3

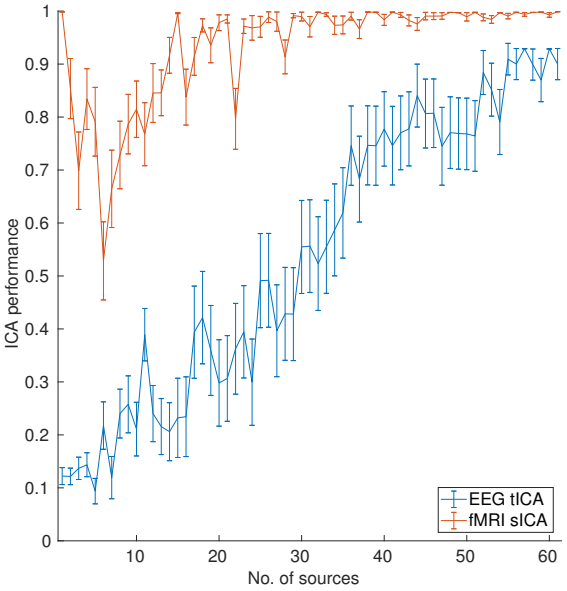

(d) Iteration 4

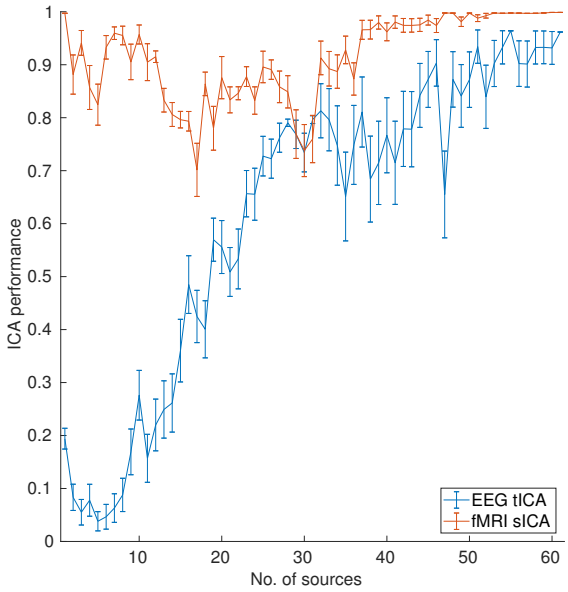

(e) Iteration 5

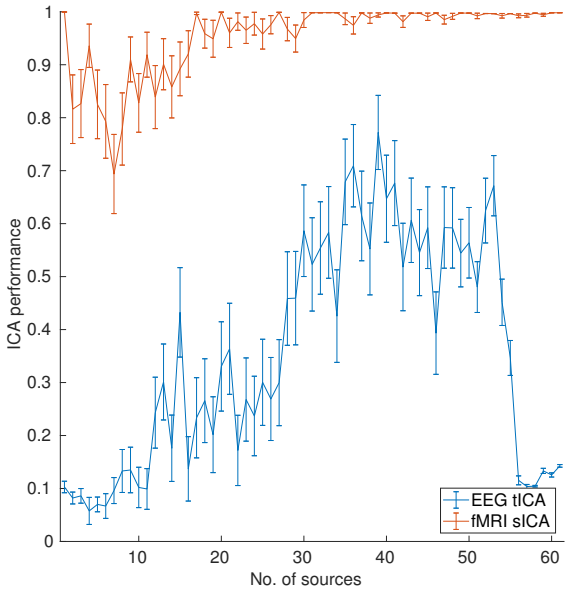

(f) Iteration 6

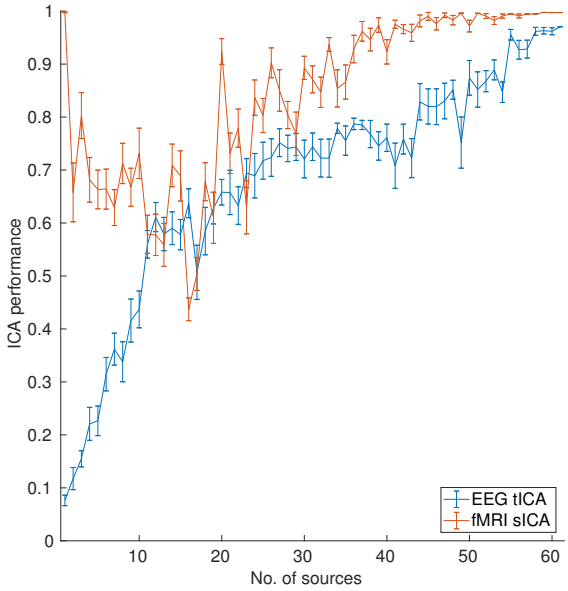

(g) Iteration 7

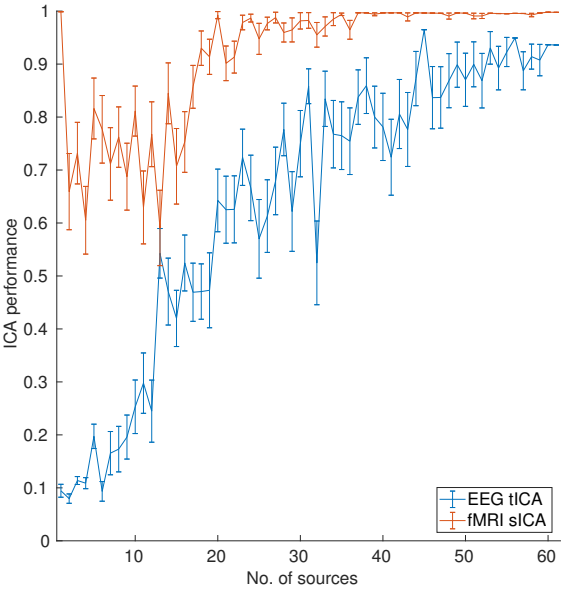

(h) Iteration 8

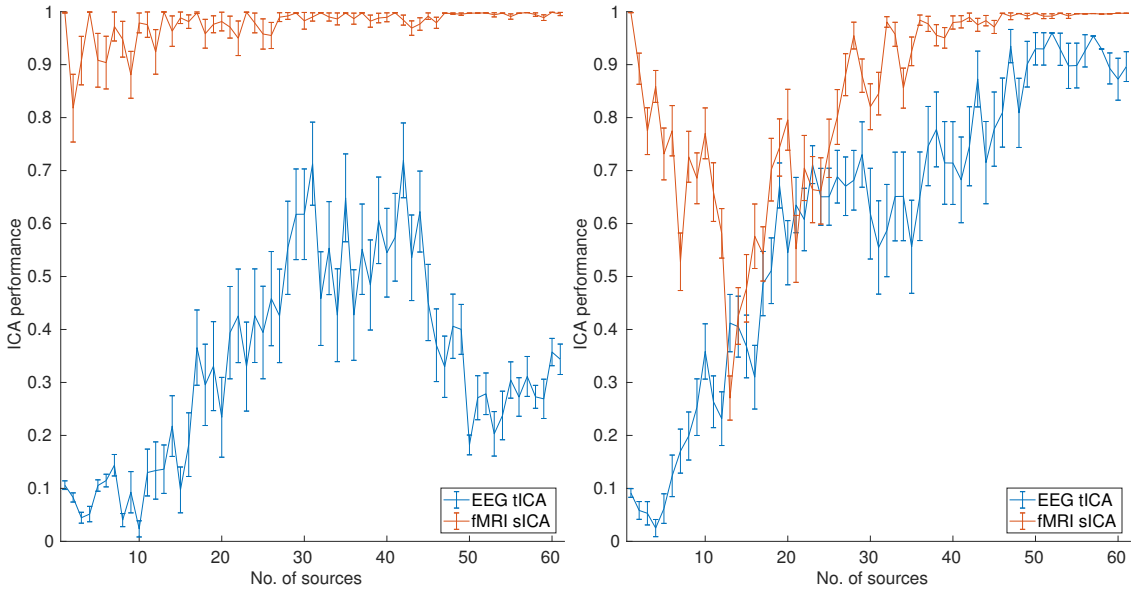

(i) Iteration 9 (j) Iteration 10

Figure A.4 ICA performance as a function of the number of sources for ten different source-to-region mappings

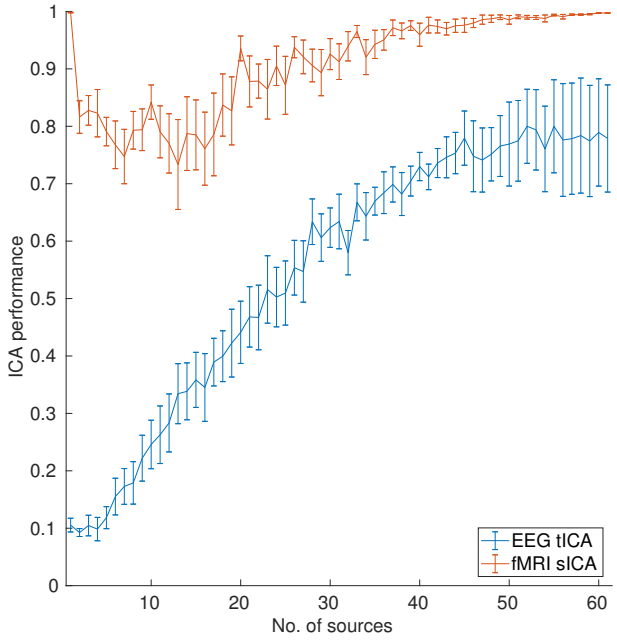

Figure A.5 Mean ICA performance as a function of the number of sources across 10 source-to-region mappings

Appendix B

Supplementary Material for Chapter 3

Why the time-course of the first source should be a linear combination of the EEG signals

Let:

$$\begin{aligned} y_1(t) &= w_{11}s_1(t) + w_{12}s_2(t) \\ y_2(t) &= w_{21}s_1(t) + w_{22}s_2(t), \end{aligned} \tag{B.1}$$

where $y_1(t)$ and $y_2(t)$ are the time-courses of two EEG channels and $s_1(t)$ and $s_2(t)$ are the time-courses of two sources. Both $y_1(t)$ and $y_2(t)$ vary because of contributions from both $s_1(t)$ and $s_2(t)$. Would a model exist for $s_1(t)$ that was a function of both $y_1(t)$ and $y_2(t)$? To answer this, Equations B.1 can be rearranged such that:

$$\begin{aligned} s_1(t) &= \frac{1}{w_{11}}[y_1(t) - w_{12}s_2(t)] \\ s_2(t) &= \frac{1}{w_{22}}[y_2(t) - w_{21}s_1(t)] \end{aligned} \tag{B.2}$$

It follows that:

$$
\begin{aligned}
s_1(t) &= \frac{1}{w_{11}}[y_1(t) - w_{12}s_2(t)] \\
&= \frac{1}{w_{11}}[y_1(t) - w_{12}\{\frac{1}{w_{22}}[y_2(t) - w_{21}s_1(t)]\}] \\
&= \frac{1}{w_{11}}[y_1(t) - \frac{w_{12}}{w_{22}}[y_2(t) - w_{21}s_1(t)]] \\
&= \frac{y_1(t)}{w_{11}} - \frac{w_{12}}{w_{11}w_{22}}[y_2(t) - w_{21}s_1(t)] \\
&= \frac{y_1(t)}{w_{11}} - \frac{w_{12}y_2(t)}{w_{11}w_{22}} + \frac{w_{12}w_{21}s_1(t)}{w_{11}w_{22}} \qquad \text{(B.3)} \\
s_1(t) - \frac{w_{12}w_{21}s_1(t)}{w_{11}w_{22}} &= \frac{y_1(t)}{w_{11}} - \frac{w_{12}y_2(t)}{w_{11}w_{22}} \\
s_1(t)[1 - \frac{w_{12}w_{21}}{w_{11}w_{22}}] &= \frac{y_1(t)}{w_{11}} - \frac{w_{12}y_2(t)}{w_{11}w_{22}} \\
s_1(t)\frac{w_{11}w_{22} - w_{12}w_{21}}{w_{11}w_{22}} &= \frac{y_1(t)}{w_{11}} - \frac{w_{12}y_2(t)}{w_{11}w_{22}} \\
s_1(t) &= (\frac{w_{11}w_{22}}{w_{11}w_{22} - w_{12}w_{21}})(\frac{y_1(t)}{w_{11}} - \frac{w_{12}y_2(t)}{w_{11}w_{22}})
\end{aligned}
$$

From Equation B.3 it is clear that, at least in the case of two sources and two EEG channels, the time-course of a source is a linear combination of the EEG signals. It is assumed that this proof can be generalised to the case of n channels and m sources.

Appendix C

Supplementary Material for Chapter 4

C.1 Comparison of source similarity domains

The source similarity matrices described in Section 4.2.6 in the main text across the four source similarity domains are shown in Figure C.1.

C.2 (p)ICA performance as a function of the number of sources

(p)ICA performance as a function of the number of sources (1 to 61, step 1) for the source similarity thresholds 0.5, 0.6, and 0.7, is shown in Figures C.2, C.3, C.4 respectively. The source similarity domain was spatiotemporal. Each column corresponds to a pICA similarity domain (temporal, spatial, spatiotemporal) and each row to a pICA similarity threshold (0.5 to 0.8, step 0.1). For each simulation, ICA was performed 30 times and the curves correspond to the mean across decompositions.

C.3 (p)ICA performance as a function of source noise dispersion

(p)ICA performance as a function of source noise dispersion (0 to 0.2, step 0.1) for the source similarity thresholds 0.5, 0.6, and 0.7, is shown in Figures C.8, C.9, C.10 respectively. The source similarity domain was spatiotemporal. Each column corresponds to a pICA similarity domain (temporal, spatial, spatiotemporal) and each row to a pICA similarity threshold (0.5

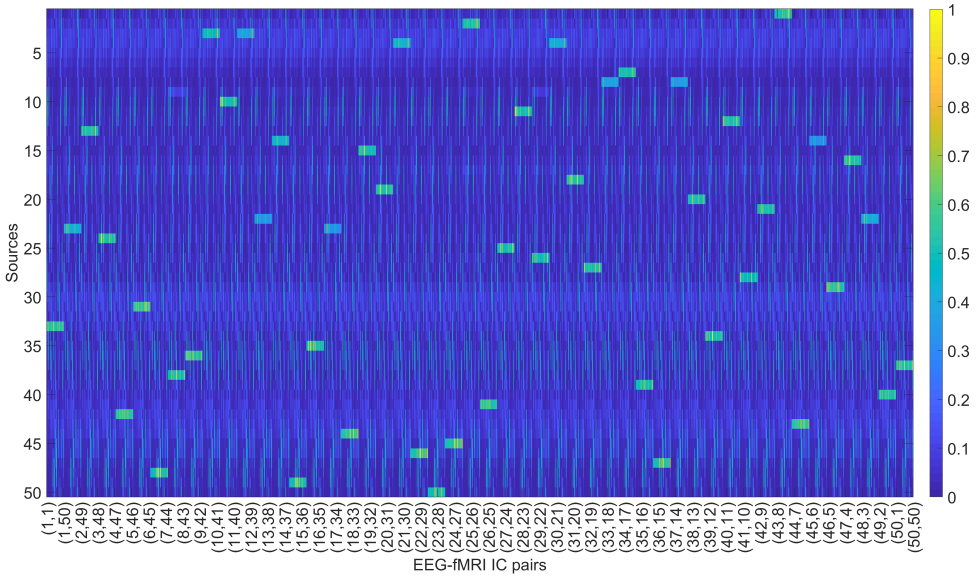

(a) Temporal

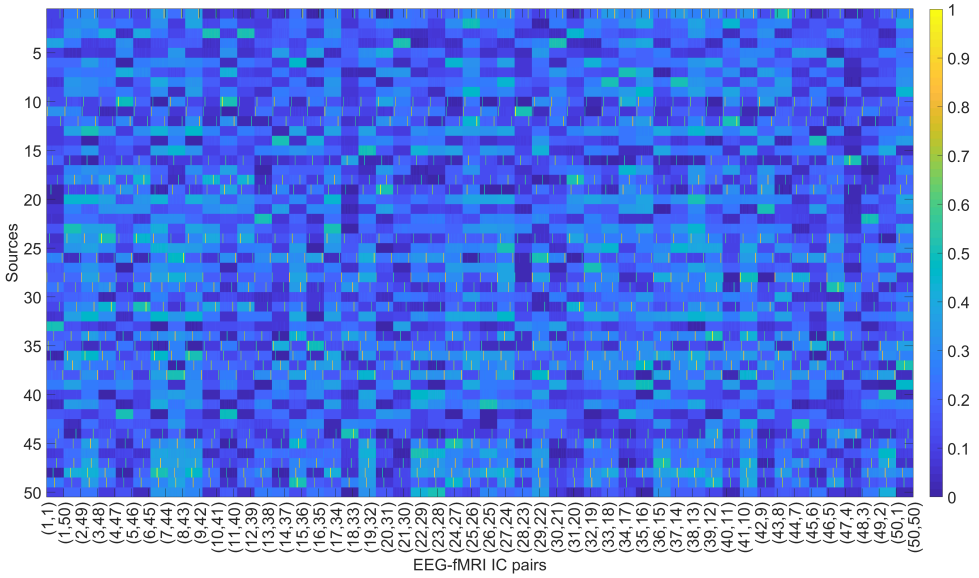

(b) Spatial

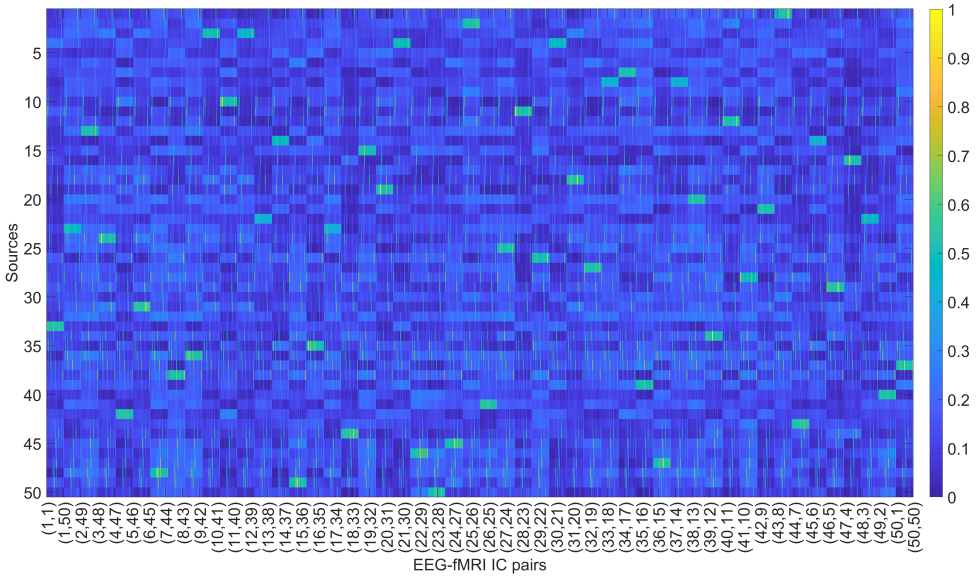

(c) Spatiotemporal

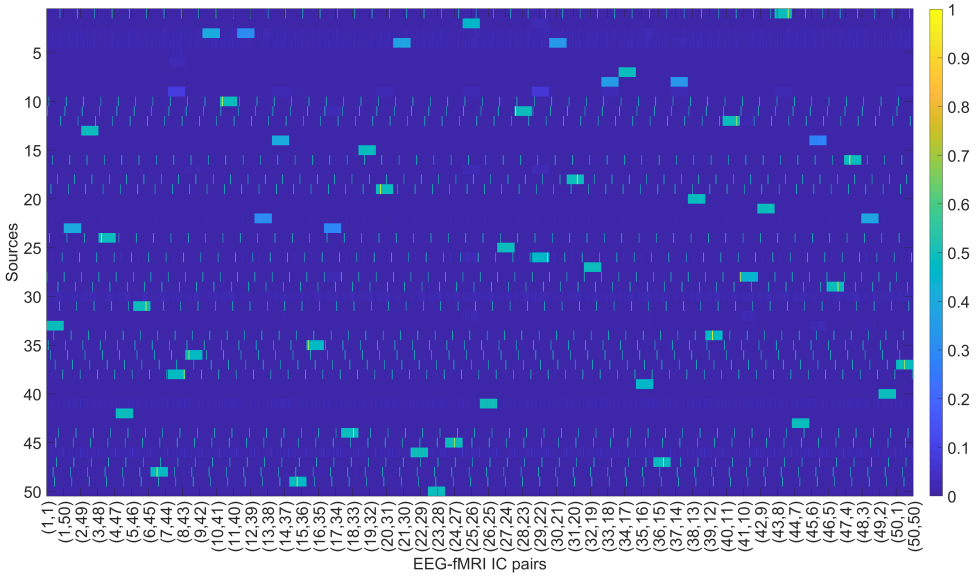

(d) Modality-specific

Figure C.1 Unthresholded similarity between all EEG-fMRI IC pairs and sources in each source similarity domain

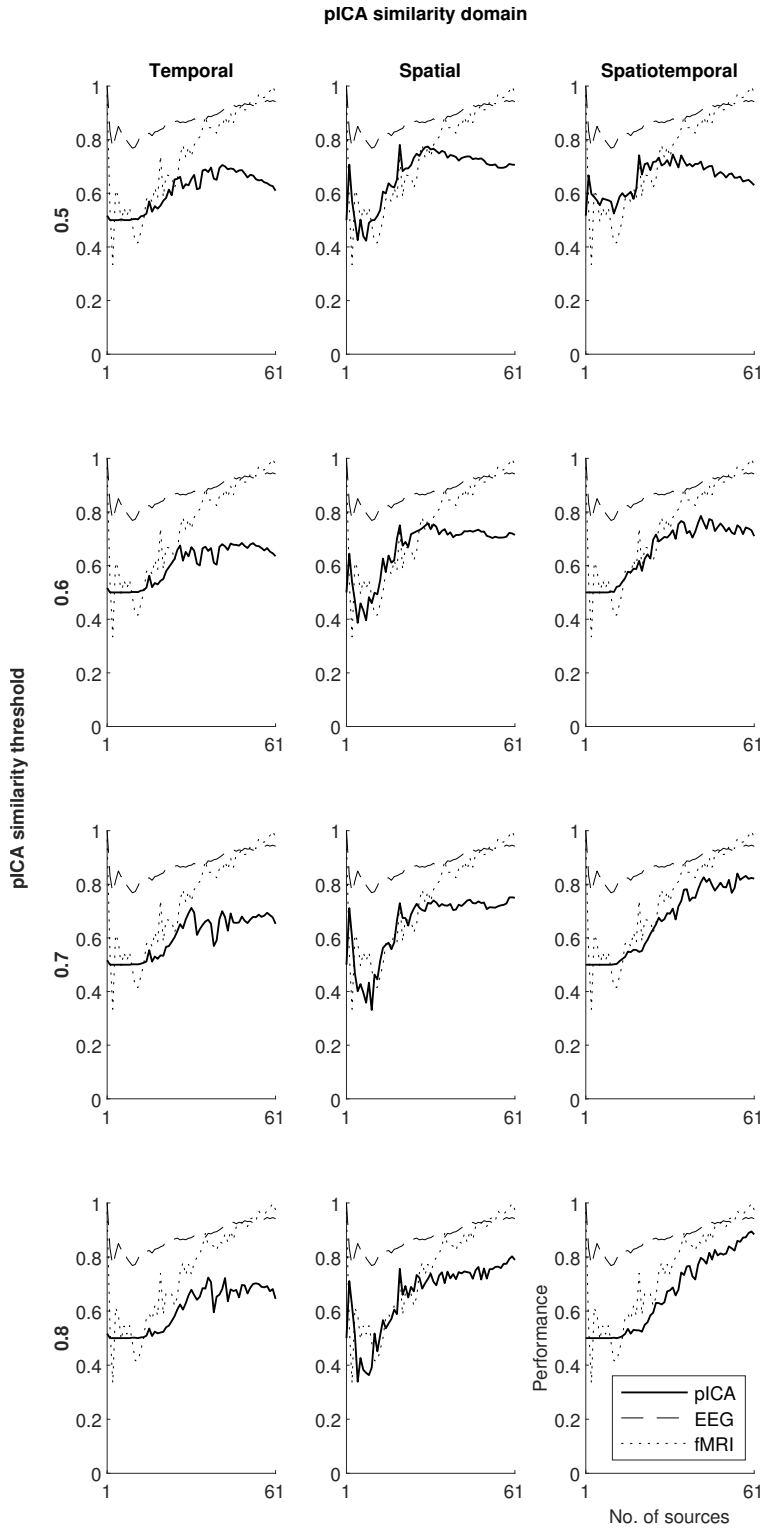

Figure C.2 (p)ICA performance as a function of the number of sources (1 to 61, step 1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.5. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0 to 0.046, EEG: 0 to 0.0325, fMRI: 0.0008 to 0.059) and are not shown to preserve visual clarity.

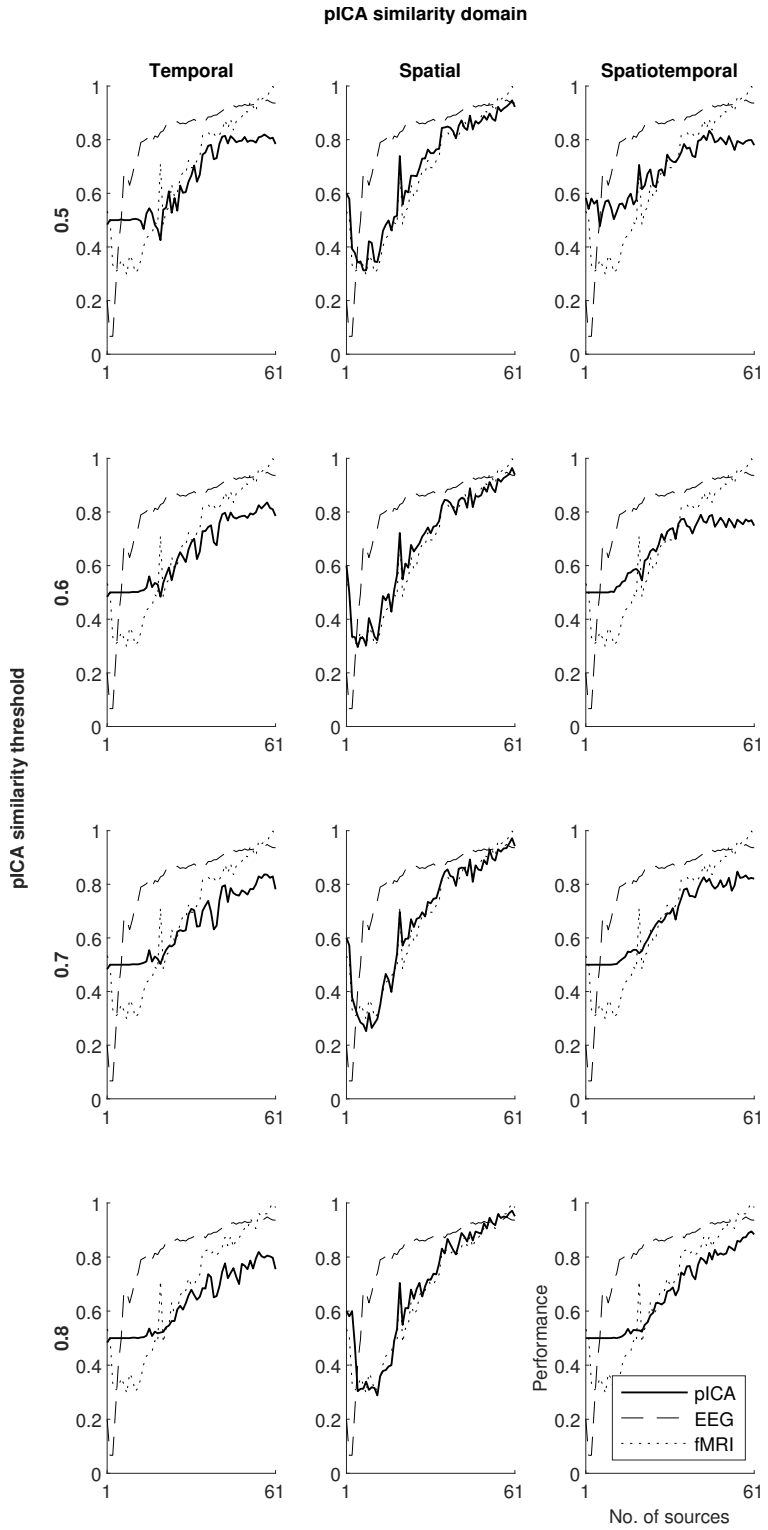

Figure C.3 (p)ICA performance as a function of the number of sources (1 to 61, step 1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.6. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0 to 0.043, EEG: 0 to 0.0325, fMRI: 0.0008 to 0.0588) and are not shown to preserve visual clarity.

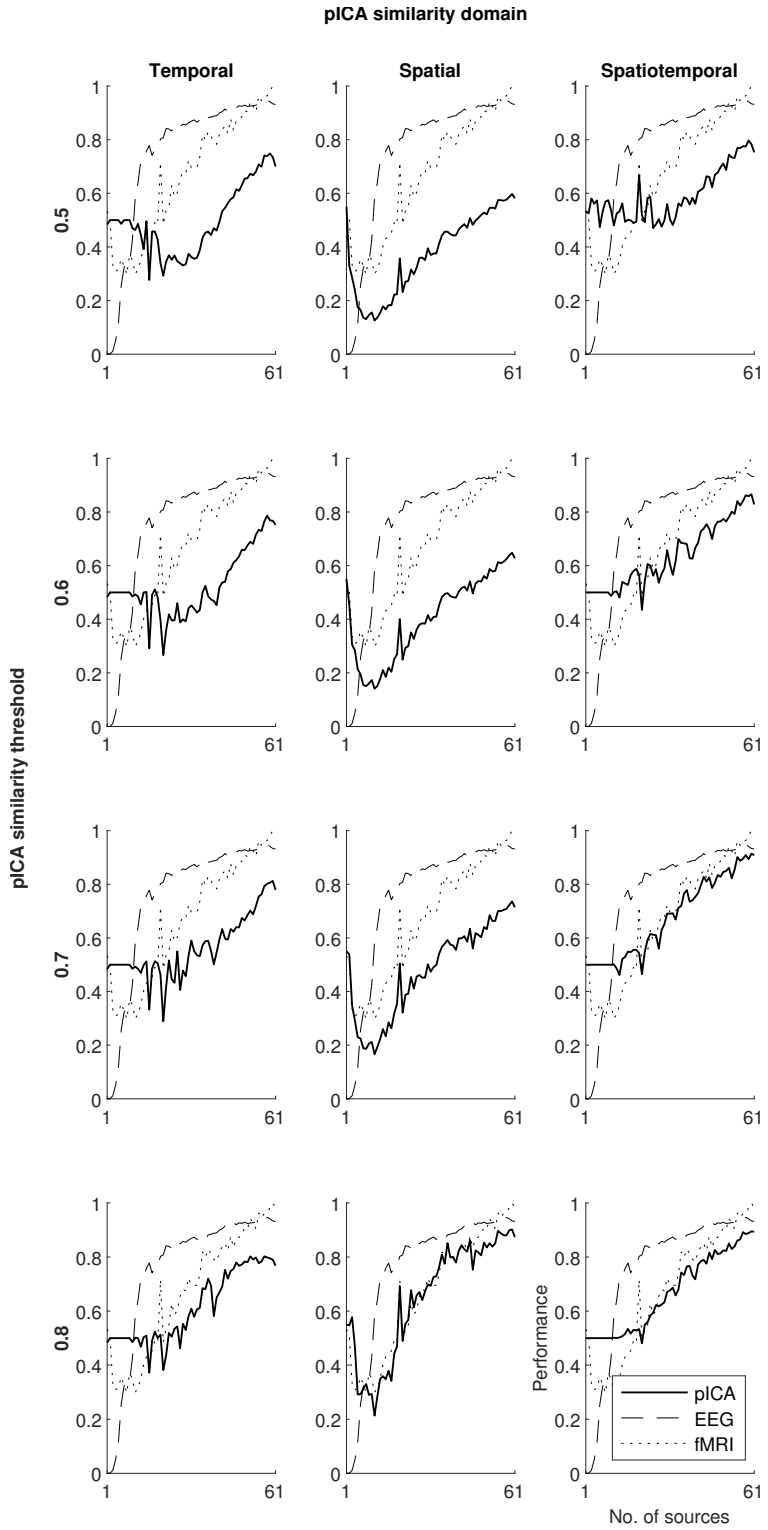

Figure C.4 (p)ICA performance as a function of the number of sources (1 to 61, step 1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.7. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0 to 0.033, EEG: 0 to 0.033, fMRI: 0.0012 to 0.059) and are not shown to preserve visual clarity.

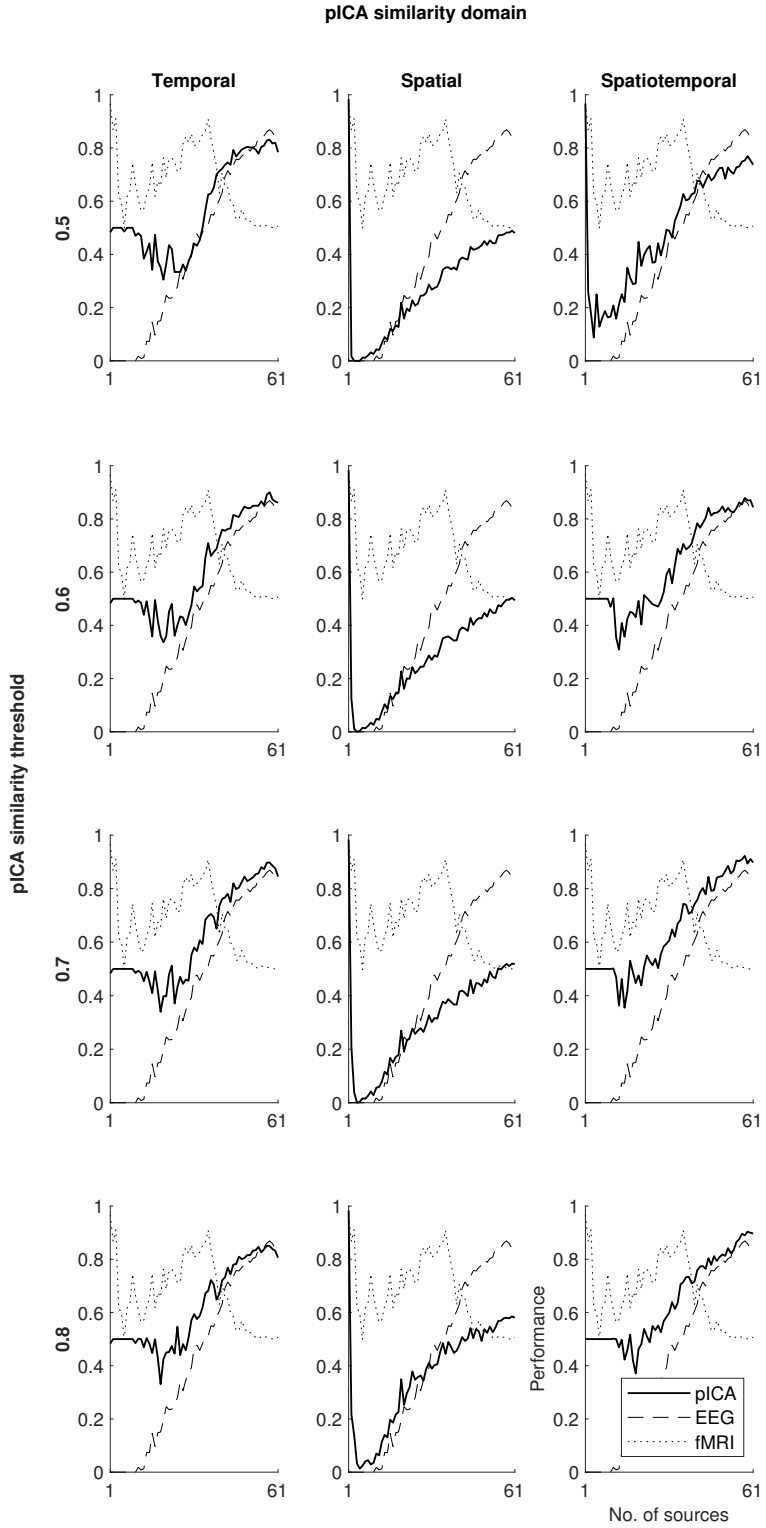

Figure C.5 (p)ICA performance as a function of the number of sources (1 to 61, step 1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was temporal and threshold was 0.8. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0 to 0.043, EEG: 0 to 0.017, fMRI: 0 to 0.033) and are not shown to preserve visual clarity.

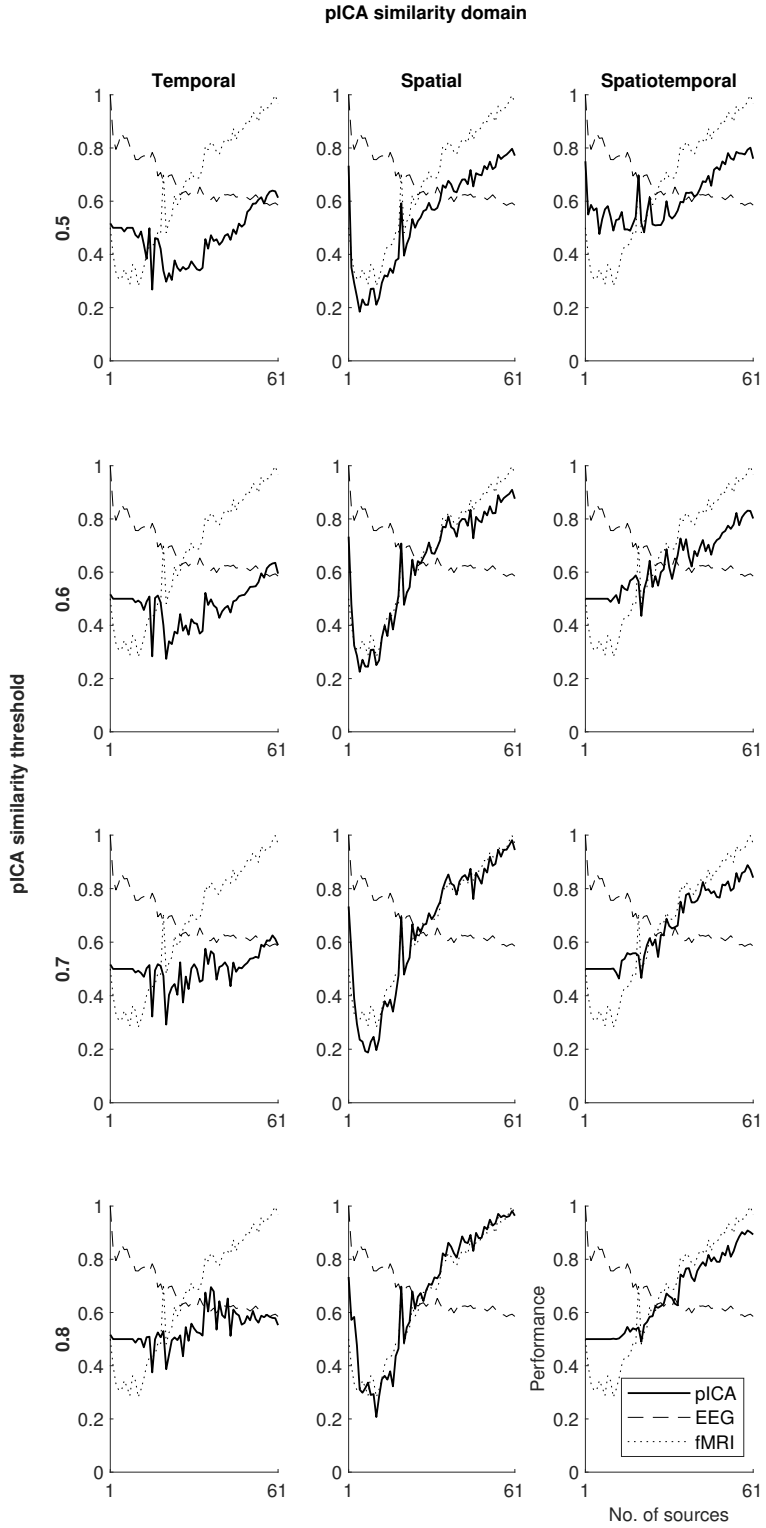

Figure C.6 (p)ICA performance as a function of the number of sources (1 to 61, step 1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatial and threshold was 0.8. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0 to 0.059, EEG: 0 to 0.028, fMRI: 0.0012 to 0.093) and are not shown to preserve visual clarity.

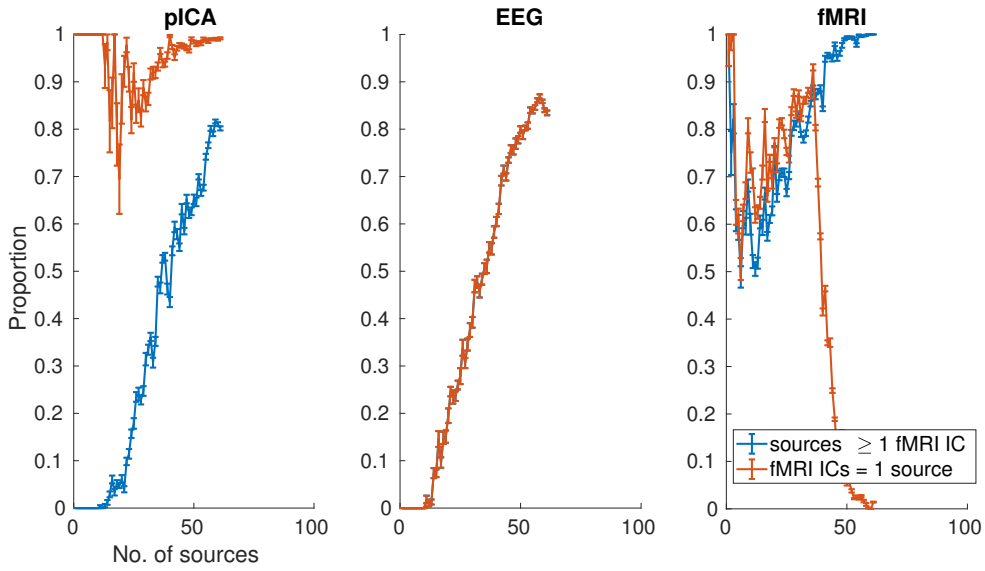

Figure C.7 The proportion of sources with at least one matching (p)IC (blue curve) and the proportion of (p)ICs matching exactly one source (orange curve) for pICA, EEG ICA, and fMRI ICA. The pICA similarity domain is spatiotemporal and the pICA similarity threshold is 0.8. The source similarity domains is temporal and the threshold is 0.8. The mean of these two quantities is used as the performance measure (see Section 4.2.5 in the main text). These quantities were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions and the errors bars to the standard error of the means.

to 0.8, step 0.1). For each simulation, ICA was performed 30 times and the curves correspond to the mean across decompositions.

C.4 (p)ICA performance as a function of source noise dispersion with dispersion varying from 0 to 2

C.4.1 Introduction

Experiment 2 was first performed with the noise dispersion varying from 0 to 2 (step 0.2) instead of from 0 to 0.2 (step 0.01). Here, the methods and results of this version of the experiment are described.

C.4.1.1 Methods

To investigate how (p)ICA performance varied as a function of source noise dispersion, the EEG-fMRI signals previously simulated in Experiment 4 of Chapter 2 were used. There were 50 sources and the first source comprised 37 regions whereas the others comprised single, non-overlapping regions. The reason the first source had 37 regions was because in Chapter 2, ICA performance was only measured for the first source. In this work, (p)ICA performance was measured across all sources (see Section 4.2.5 in the main text for a comparison of the performance measures) and, therefore, the size of the first source was not of particular interest and was kept to 37 regions simply so that previously generated data could be re-used. The noise dispersion was varied from 0 to 2 in steps of 0.2. The time-course and SNR of a single voxel is shown in Figure C.11. The hypothesis was that (p)ICA performances would decrease with noise dispersion.

C.4.2 Results

The performance of (p)ICA as a function of the source noise dispersion (0 to 2, step 0.2) is shown in Figure C.12. The source similarity domain was spatiotemporal and the source similarity threshold was 0.8. Each column corresponds to a pICA similarity domain (temporal, spatial, spatiotemporal) and each row to a pICA similarity threshold (0.5 to 0.8, step 0.1). For each simulation, ICA was performed 30 times and the curves correspond to the mean across decompositions. As the range of the standard error of the means of each curve was small (pICA: 0 to 0.023, EEG: 0 to 0.006, fMRI: 0 to 0.004), to preserve visual clarity, error bars were not shown. The figures for the other source similarity thresholds (0.5, 0.6, 0.7) are shown in Section C.4.4. The performance of pICA was greater than 0.8 only when there

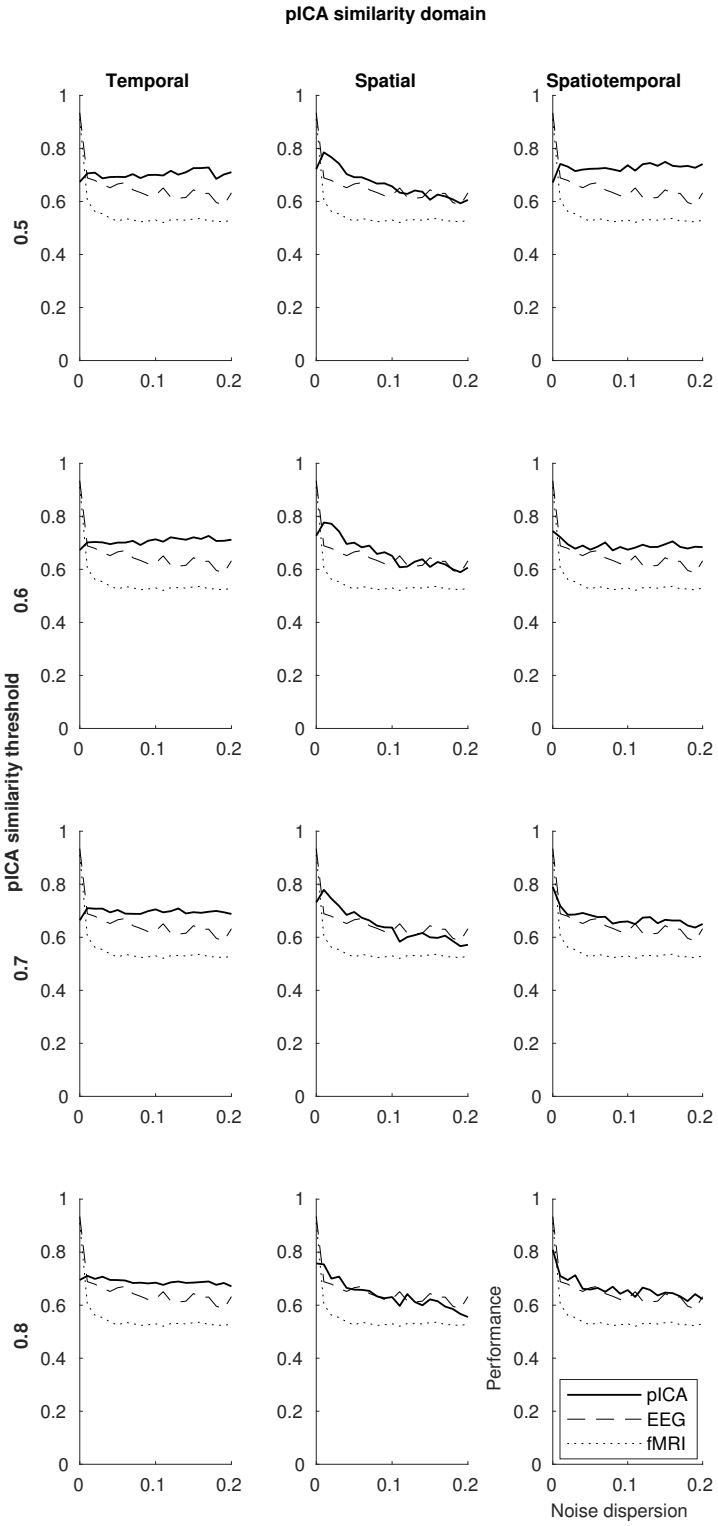

Figure C.8 (p)ICA performance as a function of the source noise dispersion (0 to 0.2, step 0.1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.5. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0.002 to 0.007, EEG: 0.003 to 0.01, fMRI: 0.002 to 0.004) and are not shown to preserve visual clarity.

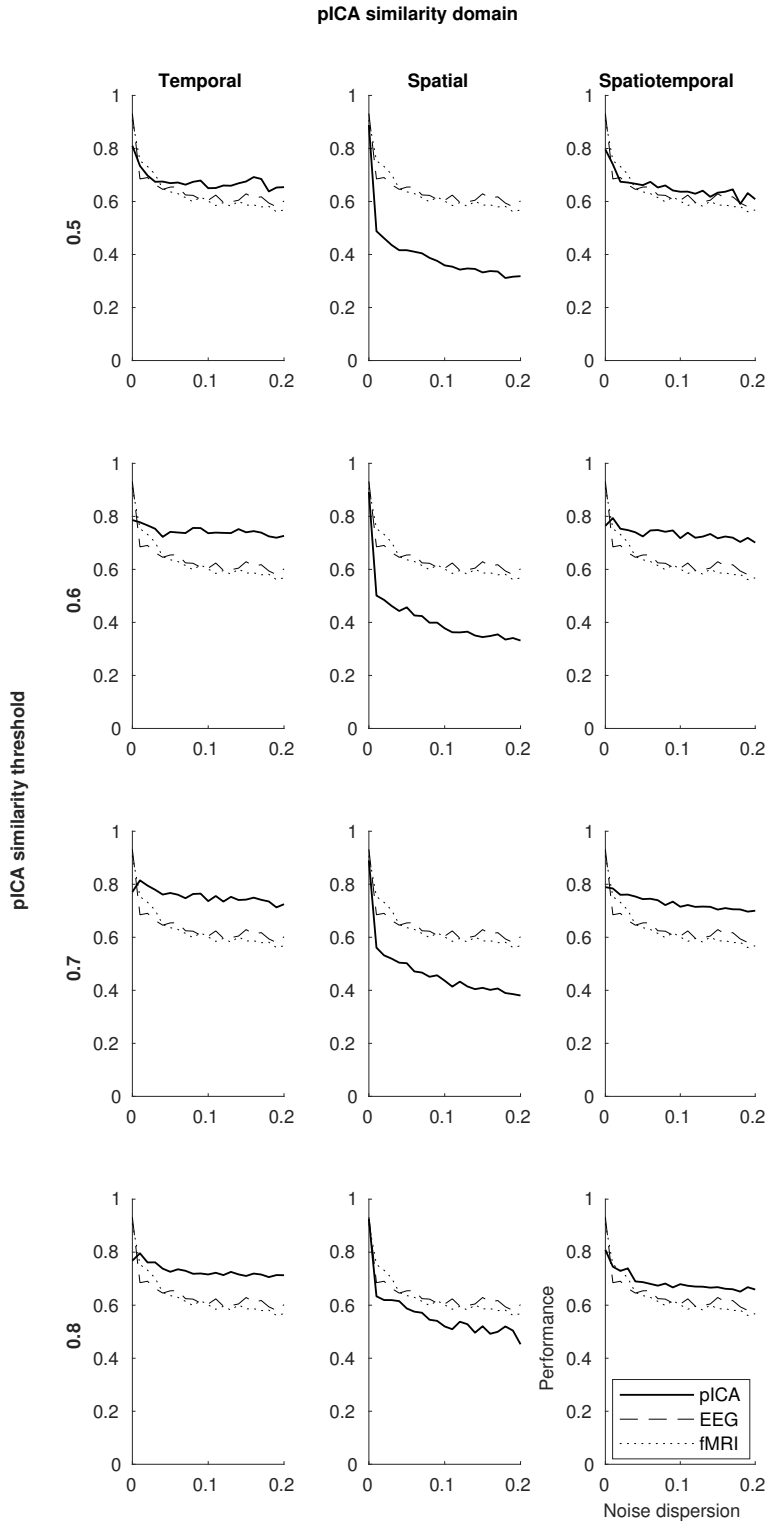

Figure C.9 (p)ICA performance as a function of the source noise dispersion (0 to 0.2, step 0.1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.6. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0.001 to 0.008, EEG: 0.003 to 0.011, fMRI: 0.002 to 0.005) and are not shown to preserve visual clarity.

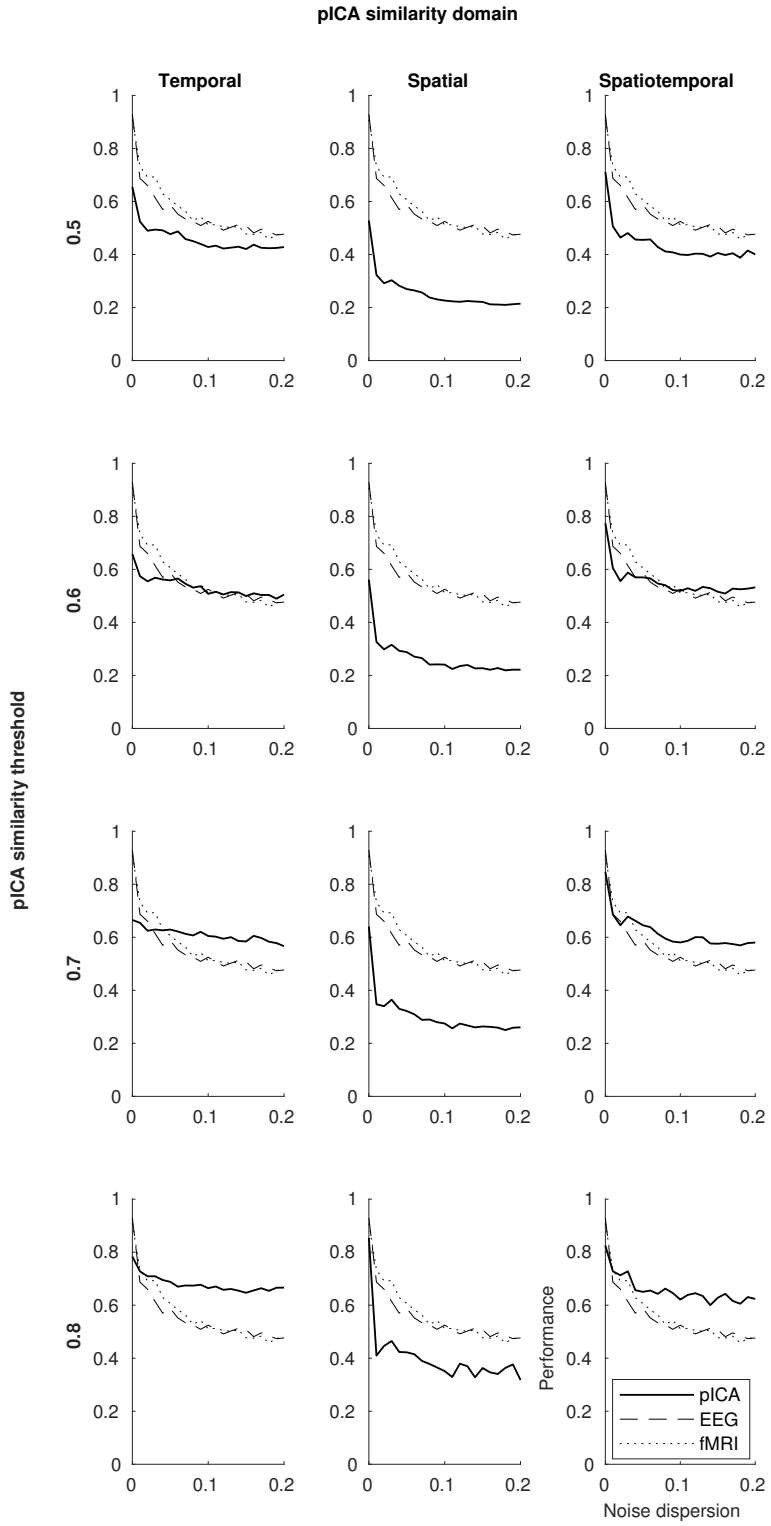

Figure C.10 (p)ICA performance as a function of the source noise dispersion (0 to 0.2, step 0.1) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.7. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0.001 to 0.08, EEG: 0.003 to 0.012, fMRI: 0.002 to 0.008) and are not shown to preserve visual clarity.

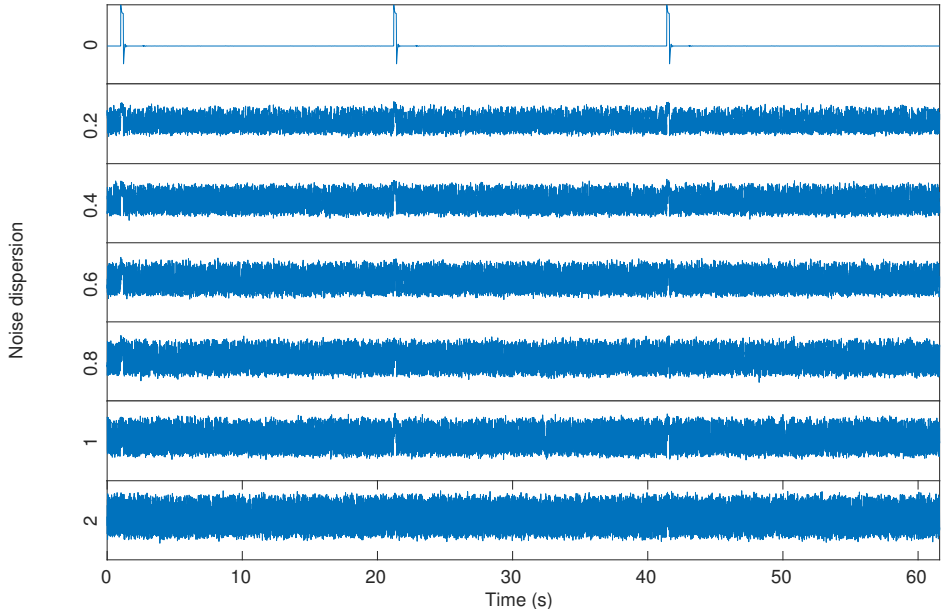

(a)

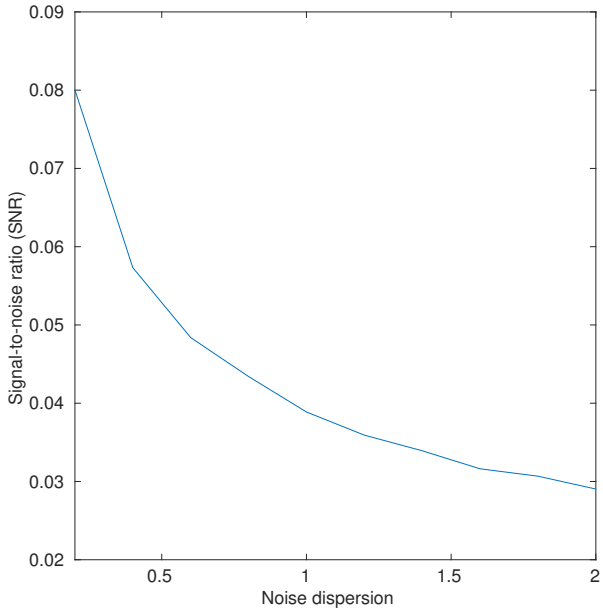

(b)

Figure C.11 The source noise dispersion was varied from 0 to 2 in steps of 0.2. For a single voxel (a) shows the local field potential (LFP) time-courses and (b) the signal-to-noise ratio (SNR) across the different noise levels (excluding noise dispersion = 0 as in this case SNR $\rightarrow \infty$).

was no noise, i.e. the dispersion was 0. In these cases, the pICA similarity domain was spatiotemporal and the threshold was 0.7 or 0.8. When the noise dispersion was 0.2, the highest performance was 0.161, and this again was in the spatiotemporal domain with a threshold of 0.8.

C.4.3 Discussion

The performance of (p)ICA decreased exponentially with noise dispersion, with pICA not performing well (> 0.8) when there was noise (dispersion > 0) in any case. When the dispersion was increased by a single step (0.2, SNR = 0.08), the highest pICA performance observed was 0.161. While this was the smallest increment possible with the data available, it was concluded that this did not suggest that pICA is not robust to a small amount of noise, as the SNR was very low. It was decided that to draw any further inference, further investigation was needed which looked at greater SNRs. Therefore, Experiment 2 (Section 4.2.8 in the main text) was performed with the dispersion varying from 0 to 0.2 in steps of 0.1.

C.4.4 Additional figures

(p)ICA performance as a function of source noise dispersion (0 to 2, step 0.2) for the source similarity thresholds 0.5, 0.6, and 0.7, is shown in Figures C.13, C.14, C.15 respectively. The source similarity domain was spatiotemporal. Each column corresponds to a pICA similarity domain (temporal, spatial, spatiotemporal) and each row to a pICA similarity threshold (0.5 to 0.8, step 0.1). For each simulation, ICA was performed 30 times and the curves correspond to the mean across decompositions.

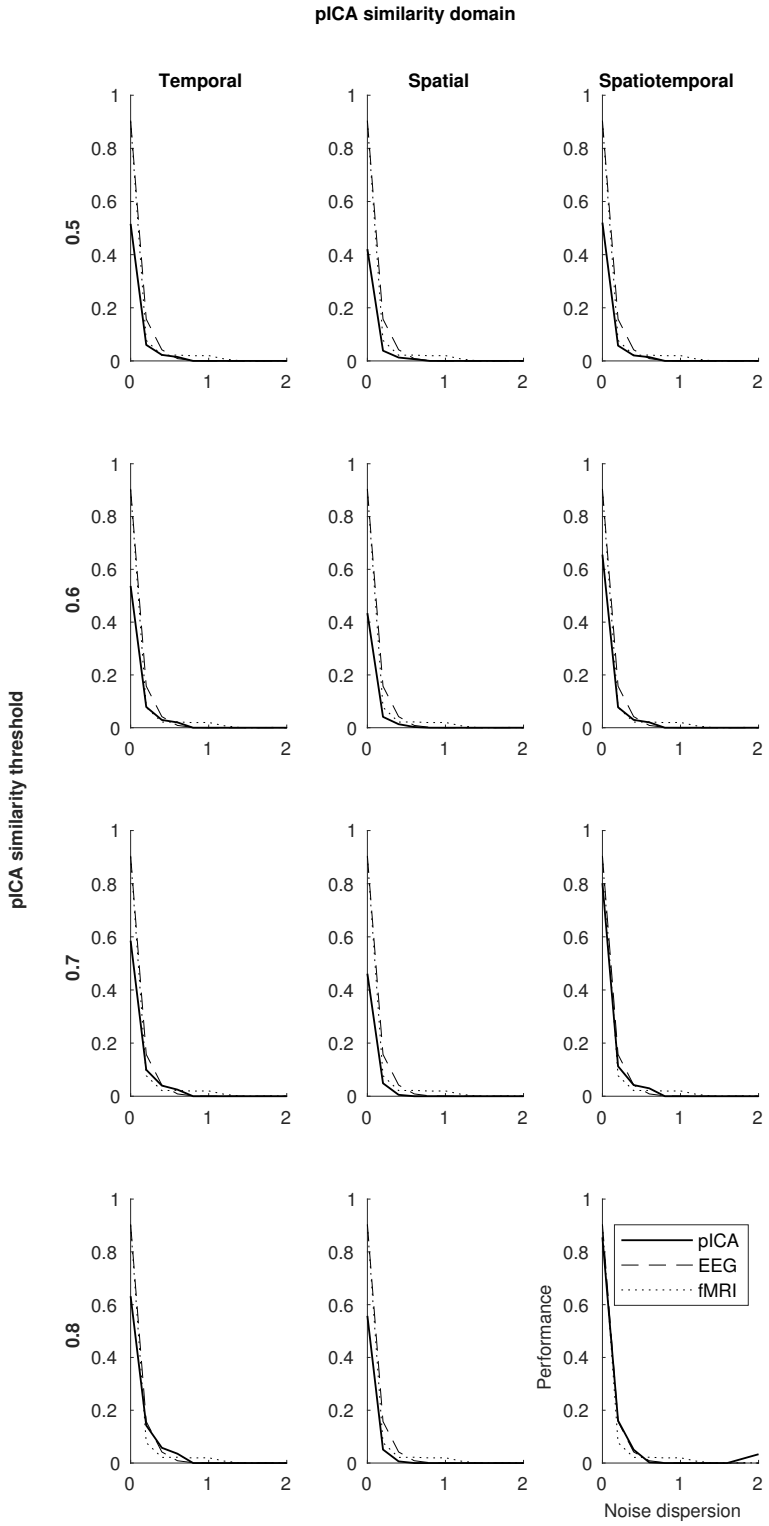

Figure C.12 (p)ICA performance as a function of the source noise dispersion (0 to 2, step 0.2) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.8. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0 to 0.023, EEG: 0 to 0.006, fMRI: 0 to 0.004) and are not shown to preserve visual clarity.

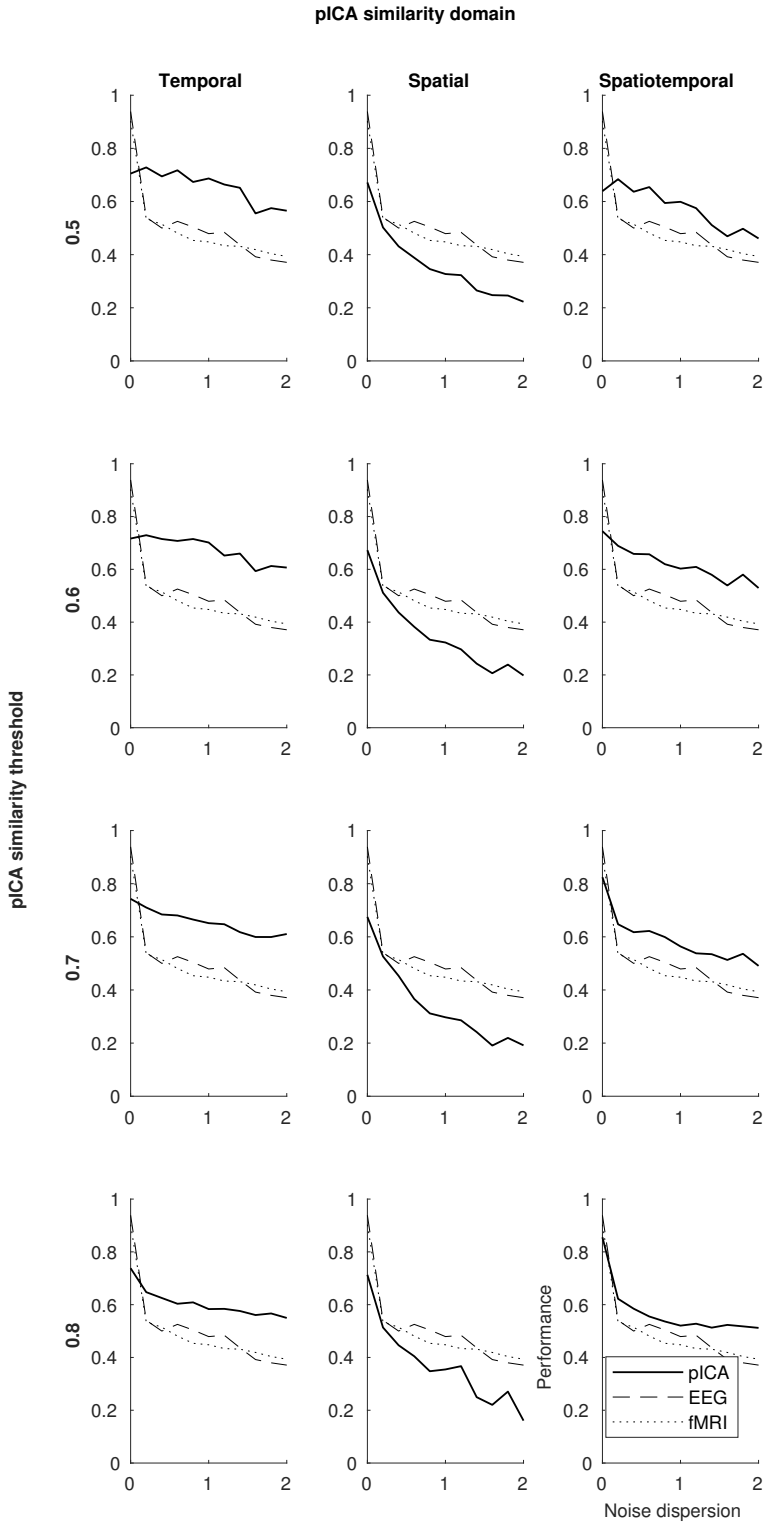

Figure C.13 (p)ICA performance as a function of the source noise dispersion (0 to 2, step 0.2) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.5. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0.001 to 0.013, EEG: 0.002 to 0.005, fMRI: 0.002 to 0.004) and are not shown to preserve visual clarity.

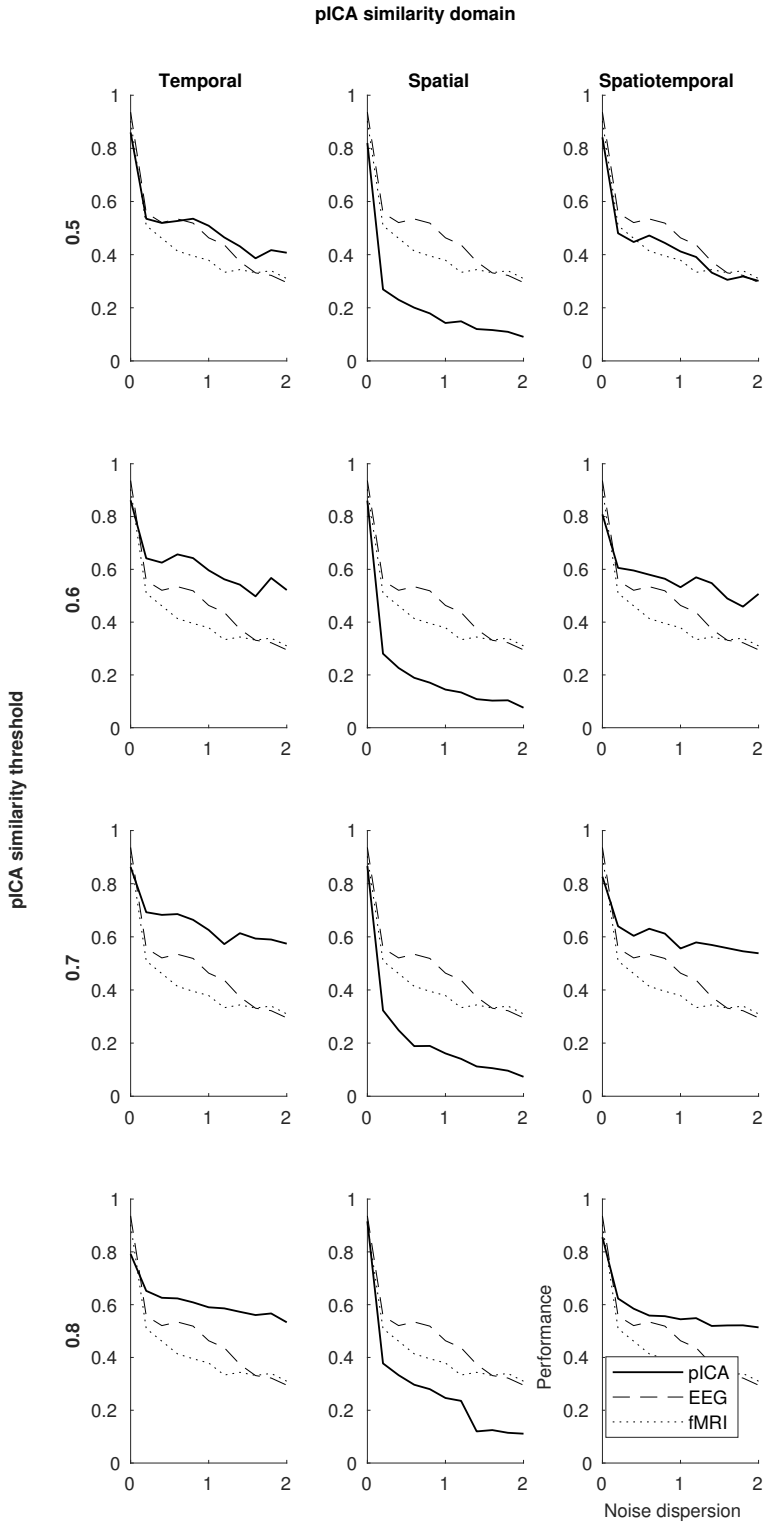

Figure C.14 (p)ICA performance as a function of the source noise dispersion (0 to 2, step 0.2) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.6. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0.001 to 0.012, EEG: 0.002 to 0.004, fMRI: 0.002 to 0.007) and are not shown to preserve visual clarity.

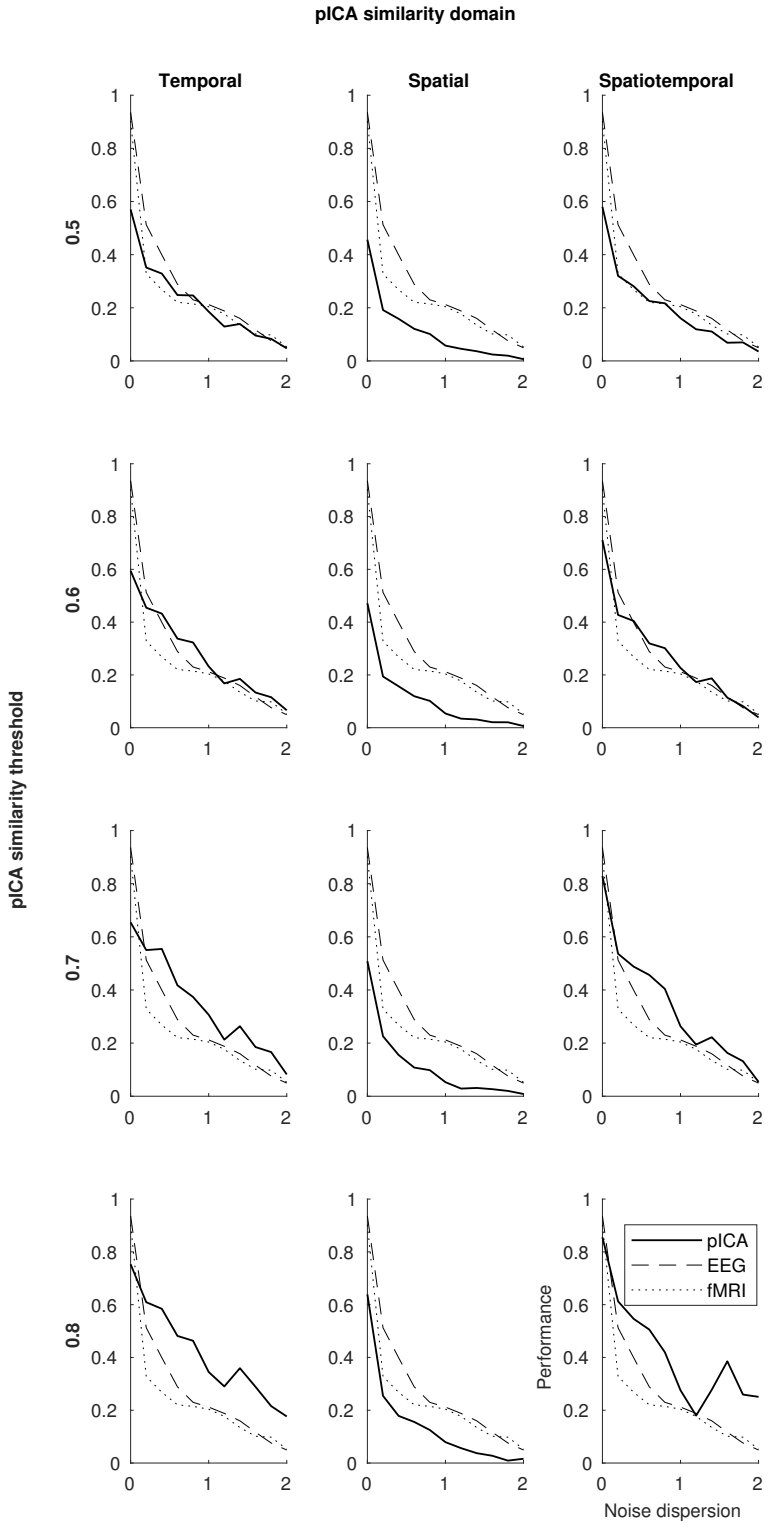

Figure C.15 (p)ICA performance as a function of the source noise dispersion (0 to 2, step 0.2) in each pICA similarity domain (columns) across pICA similarity thresholds (rows). The source similarity domain was spatiotemporal and threshold was 0.7. The performances were calculated over 30 ICA decompositions and the curves correspond to the means across the decompositions. The standard error of the means were small (pICA: 0.001 to 0.042, EEG: 0.003 to 0.006, fMRI: 0.002 to 0.007) and are not shown to preserve visual clarity.

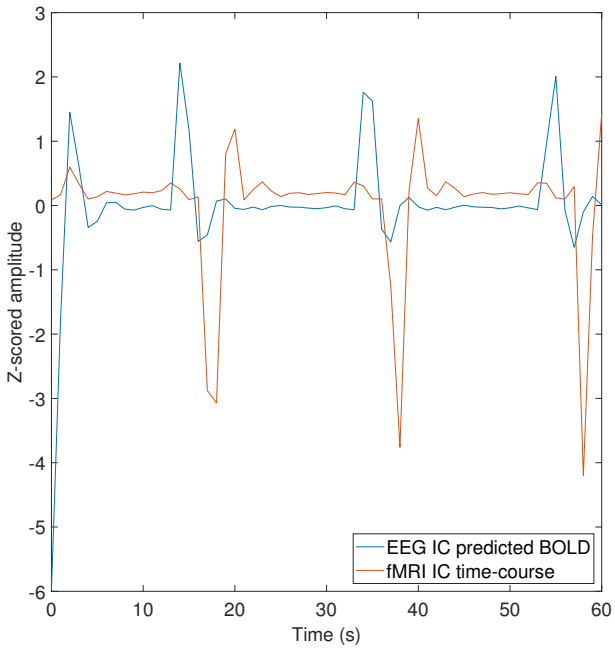

Figure C.16 The EEG (blue curve) and fMRI (orange curve) time-courses for case (b) in Figure 4.9 in the main text. In this case, the correlation between the ICs was low ($r = 0.04$) and the MI between them was high (1.04 bits). The time-courses are normalised using z-scoring for visualisation purposes

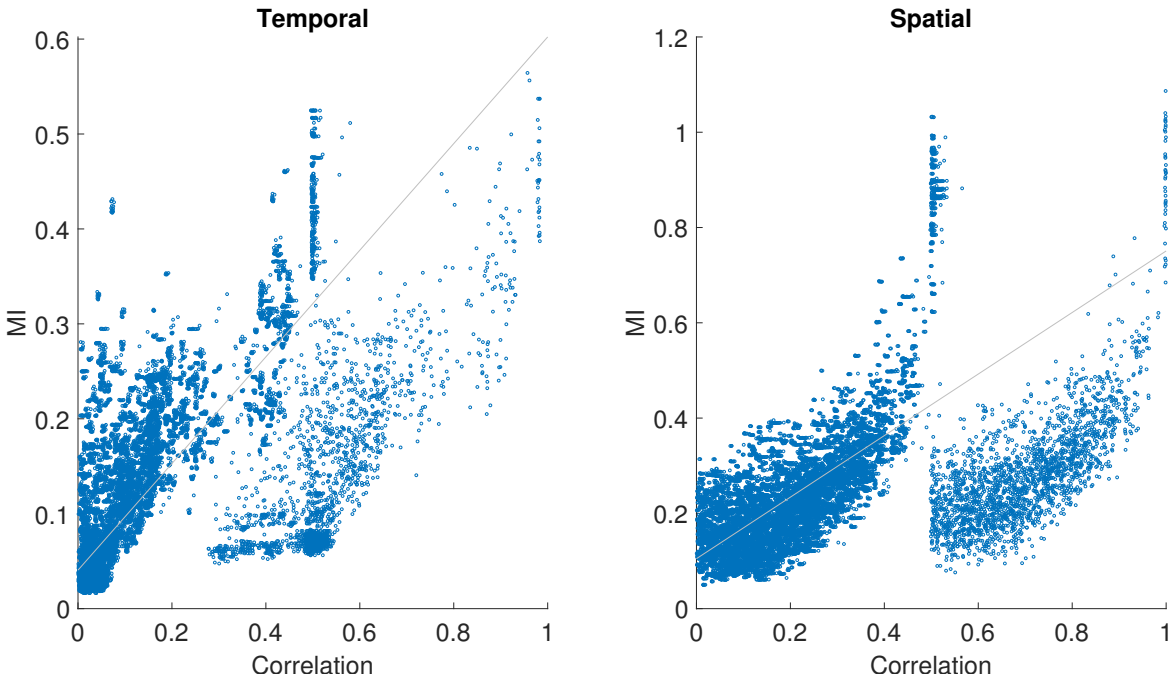

Figure C.17 Comparison of the source similarity matrices when using correlation vs. MI as the source similarity measure when the source similarity domain is temporal (left) and spatial (right). The scatterplots show the similarity between each source and each EEG-fMRI IC pair (temporal $r = 0.78$, spatial $r = 0.67$). These were obtained from a single EEG ICA and fMRI ICA decomposition on signals from a simulation with 50 orthogonal sources with no source or sensor noise.

Appendix D

Group ICA

Performing group analysis on ICA decompositions is not straightforward as ICs obtained from independent ICA decompositions are not constrained to correspond to the same sources, nor are they ranked such that IC i in two different decompositions of the same dataset corresponds to the same source. Various methods have been proposed to address this (see Erhardt et al. (2011) for review). In this appendix, we describe the temporal concatenation method proposed in Calhoun et al. (2001b)[1], as this method is used by some of the joint and parallel ICA algorithms mentioned in Section 1.5. The methods used in this thesis in the experimental chapters (2, 3, and 4), however, do not use group ICA and, therefore, this appendix does not have to be read to understand their contents. The sequence of steps involved in structuring the data to perform group ICA and obtaining subject-specific ICs and mixing matrices from the group ICs are described below.

D.1 Within-subject dimensionality reduction

Let Y_i be the $(K \times V)$ data of subject i. For example, it could be their EEG data on which temporal ICA would be performed and, therefore, the rows of Y_i would be channels and the columns time-points, or it could be their fMRI data on which spatial ICA would be performed, with the rows as time-points and the columns as voxels (see Section 1.4).

The first step that is typically applied in group ICA is to reduce the dimension of each subject's data using principal component analysis (PCA). This step is optional and is performed to reduce computational load (Calhoun et al., 2001b, 2009). Let F_i^{-1} be the

[1]Here we have largely reproduced the description of the method provided in Calhoun et al. (2009).

$(L \times K)$ dimensionality reduction matrix obtained using PCA such that

$$X_i = F_i^{-1} Y_i \tag{D.1}$$

is the $(L \times V)$ reduced data of subject i.

D.2 Concatenation

The next step is that the reduced data is concatenated across subjects, i.e. let

$$C = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_P \end{bmatrix} \tag{D.2}$$

be the $(PL \times V)$ concatenated data of all subjects, where P is the number of subjects.

D.3 Across-subject dimensionality reduction

The dimension of C is then reduced again using PCA to be equal to the number of ICs to be estimated[2], i.e. let N be the number of ICs and G^{-1} be the $(N \times PL)$ dimensionality reduction matrix such that

$$X = G^{-1} C \tag{D.3}$$

is the $(N \times V)$ reduced concatenated data across the subjects.

D.4 ICA

ICA is performed on X to obtain the $(N \times V)$ group IC matrix S and the $(N \times N)$ mixing matrix M, i.e.

$$X = MS. \tag{D.4}$$

[2]This is not necessary for FastICA (Hyvärinen and Oja, 1997), as the algorithm can estimate a non-square mixing matrix.

D.5 Back-reconstruction

The rows of S are group ICs. Often, subject-specific ICs and mixing matrices are desired, for example for drawing statistical inferences. These can be obtained using back-reconstruction. As described previously, G^{-1} is the $(N \times PL)$ group dimensionality reduction matrix. As G is the inverse of G^{-1} (assuming that it is invertible), its dimensions are $(PL \times N)$. Let G_i be the partition of G corresponding to subject i, i.e. G_1 will comprise the first L rows of G, G_2 will comprise rows $(L+1)$ to $2L$, G_i will comprise rows $[(i-1)+1]$ to $i \cdot L$. Then,

$$M_i = F_i G_i M \tag{D.5}$$

is the $(N \times N)$ subject-specific mixing matrix for subject i and

$$\begin{aligned} S_i &= M_i^{-1} Y_i \\ &= (F_i G_i M)^{-1} Y_i \end{aligned} \tag{D.6}$$

is their $(N \times V)$ subject-specific source matrix (see Calhoun et al. (2009) for proof).

D.6 Back-projection

In literature, the meaning of the term 'back-projection' is dependent on context; sometimes it is used to refer to back-reconstruction (e.g. Beckmann and Smith, 2005), and at other times to obtaining the data of an IC in a modality's signal-space (e.g. Jonmohamadi et al., 2014). In the latter case, for EEG signals, this means that the back-projection of an IC(s) is the set of EEG channel time-courses that only correspond to the activity described by that IC(s), and for fMRI signals, it is the set of voxel time-courses that only correspond to the activity described by the IC(s). A typical method for obtaining the back-projection of an IC(s) is by multiplying the the column(s) of M with the row(s) of S that correspond to the IC(s).

Appendix E

Connectivity Measures

E.1 Introduction

This appendix provides a short introduction to some functional connectivity measures that are used in EEG and fMRI analysis. These measures were considered to be an important part of this project in its early stages, when some time was spent conceiving the development of a method for EEG-fMRI analysis that looked at joint voxel-wise connectivity before it was decided to focus on parallel ICA. Then, after deciding to focus on parallel ICA, it was originally thought that we should evaluate how it performed with some of these measures. As described in Chapter 5, apart from correlation and mutual information, these measures were not used because of constraints on time. However, this appendix is considered to be a useful reference and therefore has been included in this thesis.

E.2 Notation

In this appendix, on some occasions we will define X and Y to be random variables, and in some others as stochastic processes. A stochastic process is a collection of random variables. Usually we will talk about observations of X and Y in terms of their outcomes in time. For this, we will define $x(t)$ and $y(t)$ to be signals containing values of these outcomes of X and Y respectively at time t. In most cases, X and Y will be stochastic processes, unless the measure we are discussing is explicitly defined to be between random variables, such as mutual information (see Section E.4.4).

On many occasions we will use the notation $x \in X$, where X will be a stochastic process or random variable. What we will mean by this is that x belongs to the sample space or alphabet of X. We understand that we should be using different symbols for the alphabet of a

random variable and its name, but we have decided to compromise on mathematical rigour for the sake of not having to introduce additional symbols.

E.3 Properties

E.3.1 Lower-order vs. higher-order measures

The relationship, or coupling, between signals can be linear or nonlinear. Lower-order statistics, such as correlation, can only measure the linearity of the coupling. To detect and measure nonlinear couplings, therefore, we require higher-order statistics, such as those provided by information-theoretic measures (Assecondi et al., 2015). Two examples of such measures are mutual information (see Section E.4.4) and transfer entropy (see Section E.5.5).

E.3.2 Static vs. dynamic measures

In this appendix, we make a distinction between static and dynamic coupling. Two signals have static coupling if the strength of the coupling between them is constant over time and dynamic coupling if it varies. Based on this distinction, we define static connectivity measures as those that provide a single value for the synchronisation between two signals, and dynamic connectivity measures as those that provide a time-series. In most cases, static measures can be used dynamically by using sliding-windows (Hutchison et al., 2013).

E.3.3 Functional vs. effective connectivity measures

Functional connectivity measures reflect the statistical dependence or covariation between signals (Friston et al., 2013; Friston, 2011; Vicente et al., 2011), e.g. correlation and mutual information, and they can be both directed and undirected (Friston et al., 2013). If they are undirected, they typically measure the instantaneous interactions the signals have with each other. If they are directed, they may measure time-delayed coupling or, in some cases, causal relationships in terms of Wiener causality[1].

[1]Nobert Wiener proposed a definition of causality (Wiener, 1956) that is popularly used in time-series analysis. The basic idea is that if the future of time-series Y is better predicted by a model that uses information from the history of another time-series X compared to a model that only uses information from the history of Y, then the system underlying X has some causal influence on the system underlying Y. It must be acknowledged that we have found Wiener's lecture (Wiener, 1956) in which he defined this difficult to understand because of its mathematical complexity and this explanation is based on interpretations provided by other authors (Lindner et al., 2011; Materassi, 2014; Seth et al., 2015). Granger causality (see Section E.5.6) and transfer entropy (see Section E.5.5) are causal measures based on this definition.

Effective connectivity measures reflect the influence one signal has on the other (Friston et al., 2013) and, therefore, they are directed. What separates effective connectivity measures from directed functional connectivity measures is that the latter are based on statistical dependencies of observations (signals), whereas the former are parameters in models that describe the observations, e.g. the coupling parameters in dynamic causal modelling (Daunizeau et al., 2011; Friston et al., 2003, 2013).

This distinction between functional and effective connectivity has been described in detail in Friston (2011). There appears, however, to be some misunderstanding between researchers about this classification; for example, transfer entropy is referred to as an effective connectivity measure in Vicente et al. (2011) and a directed functional connectivity measure in Friston et al. (2013). In this appendix, we stick to the definitions described above and, therefore, all the connectivity measures described in this appendix are functional connectivity measures.

E.4 Information theory

E.4.1 Introduction

In this section, we provide an introduction to a few concepts from information theory. Specifically, we describe Shannon's entropy (Shannon, 1948) (from hereon referred to as entropy), Kullback-Leibler divergence, and mutual information.

E.4.2 Entropy

Let X be a random variable. One way to define the information we get from observing an outcome of X is $-\log p_X(x)$, where x is an outcome of X and $p_X(\cdot)$ is the probability mass function of X. Intuitively, the information we get from an outcome is inversely related to its probability, i.e the occurrence of likely events provides us less information than the occurrence of unlikely ones. For example, consider tossing a fair coin. As the probability of the coin landing on a 'heads' is equal to the probability of it landing on a 'tails', either outcome provides us the same amount of information. On the other hand, if the coin had 'heads' on both sides, we would get no information by observing its outcomes, as we could predict with certainty that it would always land on a 'heads'. The entropy of X, $H(X)$, is the expected information we get from its outcomes. It is defined as:

$$H(X) = -\sum_{x \in X} p_X(x) \log p_X(x). \tag{E.1}$$

The entropy of the outcomes of a fair coin, therefore, is greater than that of the outcomes of a coin with 'heads' on both sides, which, in fact, has an entropy of 0.

E.4.3 Kullback-Leibler divergence

Let X be a random variable for which we have two probability mass functions, $p_X(\cdot)$ and $q_X(\cdot)$. Kullback-Leibler divergence is a means for quantifying the distance between these two functions. It is defined as:

$$D(p_X;q_X) = \sum_{x\in X} p_X(x)\log\frac{p_X(x)}{q_X(x)}, \tag{E.2}$$

where x is an outcome of X. In the trivial case when $p_X = q_X$, it is clear that $D(p_X;q_X) = 0$. In all other cases, $D(p_X;q_X) > 0$ (Cover and Thomas, 2006).

E.4.4 Mutual information

Let X and Y be random variables. The mutual information of X and Y, $I(X;Y)$, is a measure of the statistical dependence[2] between them. It is defined as the Kullback-Leibler divergence between the two joint probability mass functions, $p_{XY}(x,y)$ and $p_X(x)p_Y(y)$:

$$\begin{aligned} I(X;Y) &= D(p_{XY};p_X\cdot p_Y) \\ &= \sum_{x\in X}\sum_{y\in Y} p_{XY}(x,y)\log\frac{p_{XY}(x,y)}{p_X(x)p_Y(y)}, \end{aligned} \tag{E.3}$$

where x and y are outcomes of X and Y respectively, $p_X(\cdot)$ and $p_Y(\cdot)$ are the marginal probability mass functions of X and Y respectively, and $p_{XY}(\cdot)$ is their joint probability mass function. It can be seen that if X and Y are independent, then $p_{XY}(x,y) = p_X(x)p_Y(y)$ and, therefore, $I(X;Y) = 0$.

E.4.5 Estimation

Entropy and mutual information are calculated from probability density or mass functions. For real-world continuous random variables, we have to estimate these functions, as calculating the probability of each outcome by counting is typically computationally infeasible.

[2] X and Y are statistically independent when $p_{XY}(x,y) = p_X(x)p_Y(y)$, where $p_X(\cdot)$ and $p_Y(\cdot)$ are the marginal probability mass functions of X and Y respectively, and $p_{XY}(\cdot)$ is their joint probability mass function.

Therefore, for these variables, we estimate entropy and mutual information based on estimates of their probability densities or masses.

E.4.5.1 Entropy

A simple method to estimate a probability density or mass function is by grouping observations into bins (Schürmann, 2004). For example, let X be a random variable that has outcomes that are grouped into m bins having width h, n be the total number of observations, and n_i be the number of outcomes in the i-th bin. The probability that a randomly selected outcome lies in the i-th bin therefore is $p_i = n_i/n$ and the probability of an outcome within a bin is $1/h$. Put together, the estimated probability of an outcome x therefore is:

$$\begin{aligned}\hat{p}_X(x) &= \sum_{i=1}^{m}(p_i)(\frac{1}{h})1_i(x) \\ &= \frac{1}{nh}\sum_{i=1}^{m} n_i 1_i(x),\end{aligned} \tag{E.4}$$

where $1_i(x) = 1$ if x belongs to the i-th bin and 0 otherwise. Based on such an estimate of the probability mass function of X, entropy can be estimated as:

$$\hat{H}(X) = -\sum_{x\in X} \hat{p}_X(x)\log\hat{p}_X(x). \tag{E.5}$$

$\hat{H}(X)$ will underestimate the entropy of X because, within a bin, we are assuming the outcomes have a uniform distribution. By doing so, intuitively speaking, we reduce the amount of information we get from X compared to if we had been able to estimate the probability of each outcome individually. Generally, when we estimate the probability density or mass function of a random variable from a set of observations, our estimate is biased based on our sample size (our number of observations) and response quantisation (how we group outcomes into bins) (Assecondi et al., 2015).

To reduce such biases, more complex estimates of entropy, such as the Kozachenko-Leonenko estimate (Kozachenko and Leonenko, 1987), which uses nearest-neighbour techniques, are commonly used (Kraskov et al., 2004).

E.4.5.2 Mutual information

Here, a method proposed by Kraskov et al. (2004) for estimating mutual information using nearest-neighbours is described. Let the signals $x(t)$ and $y(t)$ contain observations of the

random variables X and Y respectively. At a given t, let the delay reconstruction of X, $x_d(t)$, be defined as:

$$x_d(t) = (x(t), x(t-\tau), \ldots, x(t-(dim-1)\tau)), \tag{E.6}$$

where the embedding dimension $dim \in \mathbb{Z}^+$ and the delay time $\tau \in \mathbb{Z}_{\geq 0}$. Similarly, at a given t, let the delay construction of Y, $y_d(t)$, be defined as:

$$y_d(t) = (y(t), y(t-\tau), \ldots, y(t-(dim-1)\tau)). \tag{E.7}$$

Let S be defined as the set containing paired observations of the delay reconstructed X and Y. That is, assuming $t \in \{1, 2, \ldots, n\}$, S is:

$$S = \{(x_d(1), y_d(1)), (x_d(2), y_d(2)), \ldots, (x_d(n), y_d(n))\}. \tag{E.8}$$

Let $s(t) \in S$ such that $s(t) = (x_d(t), y_d(t))$ and $s(t_k)$ be defined as the k-th nearest-neighbour of $s(t)$ when there are exactly $k-1$ points between them in S. Then, let $d_k(t)$ be a measure of the 'maximum' distance between $s(t)$ and $s(t_k)$:

$$\begin{aligned} d_k(t) &= \|s(t) - s(t_k)\| \\ &= \|(x_d(t), y_d(t)) - (x_d(t_k), y_d(t_k))\| \\ &= \max\{|x_d(t) - x_d(t_k)|, |y_d(t) - y_d(t_k)|\}. \end{aligned} \tag{E.9}$$

Mutual information, $I(X;Y)$, can then be estimated as:

$$I(X;Y) = \psi(k) + \psi(n) - \langle \psi(n_x(t)+1) + \psi(n_y(t)+1) \rangle, \tag{E.10}$$

where n is the number of observations of X and Y, $\langle \cdot \rangle$ is the mean over t, $n_x(t)$ and $n_y(t)$ are the number of neighbours $x(t)$ and $y(t)$ respectively have at t in their embedded state-spaces within $d_k(t)$, and $\psi(\cdot)$ is the digamma function.

E.5 Measures

E.5.1 Correlation

Correlation is a static, undirected measure for determining the strength of the linear relationship between two stochastic processes that have paired observations. The correlation between the stochastic processes X and Y is defined as:

$$r = \frac{\sigma_{XY}}{\sigma_X \sigma_Y}, \tag{E.11}$$

where σ_{XY} is the covariance between X and Y, and σ_X and σ_Y are their respective standard deviations.

Correlation is normalised such that $-1 \leq r \leq 1$[3] If $r = 0$, X and Y have no linear relationship. If $r > 0$, X and Y increase and decrease in the same direction, whereas if $r < 0$, they move in opposite directions. Squared correlation, r^2, indicates the proportion of the variance of X that can be explained by the variance of Y, and vice versa.

E.5.2 Cross-correlation

Cross-correlation is the correlation between two stochastic processes when one of them has been shifted in time by a delay (or lag) τ; that is, if X and Y are two stochastic processes, and $x(t)$ and $y(t)$ are their respective outcomes at time t, the cross-correlation between them at a delay τ is the correlation between $x(t)$ and $y(t+\tau)$. Cross-correlation is a directed measure, as $x(t)$ and $y(t+\tau)$ being strongly correlated may be interpreted as X influencing Y.

E.5.3 Coherence

Magnitude-squared coherence, also known simply as coherence, is calculated similarly to squared correlation. It is different from correlation in that it operates in the frequency domain. For a given frequency f, it is defined as:

$$C(f) = \frac{P_{XY}^2(f)}{P_X(f) P_Y(f)}, \tag{E.12}$$

where X and Y are two signals, $P_X(\cdot)$ and $P_Y(\cdot)$ are their respective power spectrums, and $P_{XY}(\cdot)$ is their cross-power spectrum[4] (Ashby, 2011). Coherence is normalised such that $0 \leq C(\cdot) \leq 1$.

E.5.4 Mutual information

Mutual information is defined earlier in this appendix in Section E.4.4. It has been mentioned again in this section in order to maintain the completeness of this section as a list of measures.

[3] In signal processing, correlation does not have to be normalised. In this appendix, when we refer to correlation, we refer to its normalised version, unless stated otherwise.

[4] Briefly put, $P_X(\cdot)$ is the Fourier transform of the autocorrelation (self cross-correlation) function of X, and $P_{XY}(\cdot)$ is the Fourier transform of the cross-correlation function of X and Y (for further description see Chapter 8 in Ashby (2011)).

As described earlier, mutual information is a static, undirected connectivity measure that can detect higher-order interactions. As it is a measurement of Kullback-Leibler divergence, mutual information is always non-negative.

Time-delayed mutual information is a lagged variant of this measure, analogous to how cross-correlation is a lagged variant of correlation. As it is lagged, it is directed.

E.5.5 Transfer entropy

Transfer entropy (Kantz and Schreiber, 2004; Schreiber, 2000) is a directed functional connectivity measure. It is a causality measure based on Wiener's (Wiener, 1956) definition (see Section E.3.3), and it is dynamic as it provides a time-series for the interactions between two signals. As it is an information theory-based measure, transfer entropy can measure higher-order interactions.

Let X and Y be two stochastic processes. Let $x(t)$ and $y(t)$ be their respective observations at time t. Let $x^k(t)$ refer to the sequence $((x(t), x(t-1), \ldots, x(t-k+1))$, which contains $x(t)$ and its $k-1$ previous terms, and, similarly, let $y^l(t)$ refer to the sequence $(y(t), y(t-1), \ldots, y(t-l+1))$, which contains $y(t)$ and its previous $l-1$ terms. Transfer entropy measures if Y Wiener-causes X at time t by comparing $\Pr(x(t)|x^k(t-1))$ and $\Pr(x(t)|x^k(t-1), y^l(t-1))$[5]. We interpret the latter probability being higher than the former as an indication that Y Wiener-causes X at time t. The comparison between the two conditional probability masses is done through their Kullback-Leibler divergence (see Section E.4.3). The equation for transfer entropy therefore is:

$$\begin{aligned} T_{Y\to X}(t) &= D(\Pr(x(t)|x^k(t-1), y^l(t-1)); \Pr(x(t)|x^k(t-1))) \\ &= \sum \Pr(x(t), x^k(t-1), y^l(t-1)) \log \frac{\Pr(x(t)|x^k(t-1), y^l(t-1))}{\Pr(x(t)|x^k(t-1))}, \end{aligned} \tag{E.13}$$

where $\sum$ is the sum over $x(t) \in X(t)$, $x^k(t-1) \in X^k(t-1)$, and $y^l(t-1) \in Y^l(t-1)$. As it is a measurement of Kullback-Leibler divergence, transfer entropy is always non-negative.

[5] In this section, we do not specify individual probability mass functions for the sake of keeping the expressions visually tidy. Instead, we use $\Pr(\cdot)$ to be the probability of the argument calculated from the probability mass function of the implied random variable, e.g. $\Pr(x(t))$ is the probability of the stochastic process X having the outcome $x(t)$ according to the probability mass function of the random variable $X(t)$.

E.5.6 Granger causality

Granger causality (Granger, 1969) is a linear autoregressive model based directed functional connectivity measure. Like transfer entropy, it is also a causality measure based on Wiener's (Wiener, 1956) definition (see Section E.3.3). Unlike transfer entropy, it can only measure linear interactions between signals.

Let X and Y be two stochastic processes. Let $x(t)$ and $y(t)$ be their respective observations at time t. Let $x^k(t)$ refer to the sequence $((x(t), x(t-1), \ldots, x(t-k+1))$, which contains $x(t)$ and its $k-1$ previous terms, and, similarly, let $y^l(t)$ refer to the sequence $(y(t), y(t-1), \ldots, y(t-l+1))$, which contains $y(t)$ and its previous $l-1$ terms. The Granger causality of Y on X is calculated by comparing the following two linear autoregressive models:

$$x(t) = \sum_{i=1}^{k} a_i x(t-i) + \eta_X(t), \tag{E.14}$$

$$x(t) = \sum_{i=1}^{k} b_i x(t-i) + \sum_{j=1}^{l} c_j y(t-j) + \eta_{X|Y}(t), \tag{E.15}$$

where all η are drawn with replacement from the normal distribution and all $a, b, c \in \mathbb{R}$.

Given X and Y, the parameters of these models are typically estimated using least squares. The two models are then compared in terms of their error variances. Let σ_X^2 be the variance of η_X and σ_{XY}^2 be the variance of $\eta_{X|Y}$. If Y Granger-causes X, $\sigma_{X|Y}^2$ is expected to be less than σ_X^2. According to Ashby (2011), Granger causality is measured as follows:

$$F_{Y \to X} = \ln \frac{\sigma_X^2}{\sigma_{X|Y}^2}. \tag{E.16}$$

The fit of the model described in Equation (E.15) will always be as good or better than the model in Equation (E.14), because the former will contain all the regressors in the latter, and, therefore, $\sigma_{X|Y}^2 \leq \sigma_X^2$. Hence, Granger causality is always non-negative.

Appendix F

Ethics Application for Retinotopy Experiment

University of Reading

Coordinator for Quality Assurance in Research
Dr Mike Proven, BSc(Hons), PhD

Academic and Governance Services

Whiteknights House
Whiteknights, PO Box 217
Reading RG6 6AH

Dr Etienne Roesch
School of Psychology and Clinical Language Sciences
University of Reading
RG6 6AL

6 April 2016

Dear Etienne

UREC 16/20: Joint recording of EEG and fMRI-BOLD during retinotopic presentation *Favourable opinion*

Thank you for your application (original email dated 14 August 2015 from Catriona Scrivener, and including attachments, refers). I can confirm that the Chair is pleased to confirm a favourable ethical opinion on the basis of the information that was reviewed and the clarification from the School REC.

Please note that the Committee will monitor the progress of projects to which it has given favourable ethical opinion approximately one year after such agreement, and then on a regular basis until its completion.

Please also find attached Safety Note 59: Incident Reporting in Human Interventional Studies at the University of Reading, to be followed should there be an incident arising from the conduct of this research.

The University Board for Research and Innovation has also asked that recipients of favourable ethical opinions from UREC be reminded of the provisions of the University Code of Good Practice in Research. A copy is attached and further information may be obtained here: http://www.reading.ac.uk/internal/res/QualityAssuranceInResearch/reas-RSqar.aspx.

Yours sincerely

Dr M J Proven
Coordinator for Quality Assurance in Research (UREC Secretary)
cc: Dr John Wright (Chair); Professor Laurie Butler (Head of School); Catriona Scrivener (PhD student)

Figure F.1 Ethical approval letter

Application UREC.pdf Application UREC.pdf Application UREC.pdf

Research Ethics Committee

University of Reading

Project Submission Form

Note All sections of this form should be completed.
Please continue on separate sheets if necessary.

Principal Investigator: .....Etienne B Roesch.....

School: .....School of Systems Engineering.....

Title of Project: Joint recording of EEG and fMRI-BOLD during retinotopic presentation

Proposed starting date: ...03/08/2015.....

Brief description of Project:

Participants will be presented with visual stimuli, while fMRI-BOLD and EEG data will be recorded. Collectively, these measures will converge to facilitate the characterization of the neural circuit involved in visual perception, and will be used as benchmark data for the development of novel analysis methods for the joint data.

I confirm that to the best of my knowledge I have made known all information relevant to the Research Ethics Committee and I undertake to inform the Committee of any such information which subsequently becomes available whether before or after the research has begun.

I confirm that I have given due consideration to equality and diversity in the management, design and conduct of the research project.

I confirm that if this project is an interventional study, a list of names and contact details of the subjects in this project will be compiled and that this, together with a copy of the Consent Form, will be retained within the School for a minimum of five years after the date that the project is completed.

Signed:

..... Date: 30/07/15

(Investigator)

..... Date: .....

(Head of School or authorised Head of Department)

.... Date: 5.8.15

(Student -where applicable)

30/07/15

30/07/15

Application UREC.pdf Application UREC.pdf Application UREC.pdf

Checklist

1. This form is signed by my Head of School (or authorised Head of Department) ☑

2. The Consent form includes a statement to the effect that the project has been reviewed by the University Research Ethics Committee and has been given a favourable ethical opinion for conduct ☑

3. I have made, and explained within this application, arrangements for any confidential material generated by the research to be stored securely within the University and, where appropriate, subsequently disposed of securely. ☑

4. I have made arrangements for expenses to be paid to participants in the research, if any, OR, if not, I have explained why not. ☑

5. EITHER

   (a) The proposed research does not involve the taking of blood samples; ☑

   OR

   (b) For anyone whose proximity to the blood samples brings a risk of Hepatitis B, documentary evidence of immunity prior to the risk of exposure will be retained by the Head of School or authorized Head of Department. ☐

   Signed:

   ..... Date.....

   (Head of School or

   authorised Head of Department)

6. EITHER

   (a) The proposed research does not involve the storage of human tissue, as defined by the Human Tissue Act 2004; ☑

   OR

   (b) I have explained within the application how the requirements of the Human Tissue Act 2004 will be met. ☐

7. EITHER

(a) The proposed research will not generate any information about the health of participants; ☐

OR

(b) If the research could reveal adverse information regarding the health of participants, their consent to pass information on to their GP will be included in the consent form and in this circumstance I will inform the participant and their GP, providing a copy of the relevant details to each and identifying by date of birth ☑

OR

(c) I have explained within the application why (b) above is not appropriate. ☐

8. EITHER

(a) the proposed research does not involve children under the age of 5; ☑

OR

(b) My Head of School (or authorised Head of Department) has given details of the proposed research to the University's insurance officer, and the research will not proceed until I have confirmation that insurance cover is in place. ☐

Signed:

..... Date.....

(Head of School or authorised Head of Department)

This form and further relevant information (see Sections 5 (b)-(e) of the Notes for Guidance) should be returned, both electronically and in hard copy, to:

Dr Mike Proven
Coordinator for Quality Assurance in Research
Whiteknights House

You will be notified of the Committee's decision as quickly as possible, and you should not proceed with the project until a favourable ethical opinion has been passed.

Application UREC.pdf Application UREC.pdf Application UREC.pdf

School of

Application Form

SECTION 1: APPLICATION DETAILS
1.1 Project Title: Joint recording of EEG and fMRI-BOLD during retinotopic presentation Date of Submission: 31/07/2015 Proposed start date: 03/08/2015 Proposed End Date: 01/03/2015
1.2 Principal Investigator [supervisor name, if student project]: Etienne B. Roesch (Lecturer) Email address: Other applicants (role): Asad Malik Institution/Department: SSE Email address: Other applicants (role): Catriona Scrivener Institution/Department: SSE Email address:
1.3 Project Submission Declaration I confirm that to the best of my knowledge I have made known all information relevant to the Research Ethics Committee and I undertake to inform the Committee of any such information which subsequently becomes available whether before or after the research has begun. I understand that it is a legal requirement that both staff and students undergo Criminal Records Checks when in a position of trust (i.e. when working with children or vulnerable adults).

Application UREC.pdf Application UREC.pdf Application UREC.pdf

I confirm that a list of the names and addresses of the subjects in this project will be compiled and that this, together with a copy of the Consent Form, will be retained within the School for a minimum of five years after the date that the project is completed.

Signed (Principal Investigator) Date: 30/07/15

Signed (Student) Date: 30/07/15

Signed (Student) Date: 30/07/15

1.4

University Research Ethics Committee Applications

Projects expected to require review by the University Research Ethics Committee must be reviewed by the Chair of the School Ethics Committee or the Head of School before submission.

Signed..... (Chair of School Committee) Date:.....

Signed..... (Head of School) Date:.....

1.5

External research ethics committees

Please provide details below of other external research ethics committees to which this project has been submitted, or from whom approval has already been granted [e.g. NHS Committee]

Name of committee	Date of submission/approval	Reference	Status

Application UREC.pdf Application UREC.pdf Application UREC.pdf

SECTION 2: PROJECT DETAILS

2.1

Lay summary

Functional magnetic resonance imaging (fMRI) is capable of localizing neural activity with very fine resolution of the order of millimeters; yet because fMRI is a technique based on blood hemodynamics its temporal resolution is of the order of seconds. So essentially, fMRI can inform precisely where neural activity is occurring, but has little power to say when the activity is occurring. In contrast, electroencephalography (EEG) has the ability to record neural activity with millisecond accuracy, but its ability to precisely localize the source of the activity is very imprecise for a number of reasons. It can be said that EEG can inform precisely when neural activity is occurring, but has little power to describe where the activity is occurring.

Recent technological advances make it possible for scientists to record both EEG and fMRI data concurrently. By simultaneously recording these two modalities, scientists are able to precisely say **where and when** neural activity is occurring, whereas previously researchers had to settle for one or the other. This has been extremely problematic in the investigation of emotional experience, which has a floating nature, is never exactly the same and cannot be repeated. Concurrent measurement of EEG and fMRI data thus provides a unique opportunity to investigate the common neural source for both EEG and fMRI-BOLD activation.

Retinotopic stimulation, as planned in the present study, is one of the simplest kind of visual stimulation, where the participant is asked to fixate at the centre of the screen, while rotating or expanding wedges made of alternating checkerboard patterns are presented at the edge of the visual field. In the present study, participants will not be asked to respond to the stimuli, and will only view blocks of stimuli passively. This extremely simple design offers the advantage of producing extremely reproducible and robust neural activation, which can be located very precisely, in space and time. They are thus very suitable as baseline and benchmark for the development of novel analysis methods for the joint recording of EEG and fMRI-BOLD data.

The protocol for this study has been reviewed in the MR working group at the Centre for Integrative Neuroscience and Neurodynamics (CINN), as required by the MR Study Guidelines, and has been planned according to the guidelines and standard practices in place in CINN for the joint recording of EEG and fMRI.

2.2

Procedure

Up to 30 participants will take part in a passive viewing study, where they will be presented

with visual stimuli designed to elicit activation of the occipital cortex (retinotopy). The design of the experiment is one of the simplest, where the participant is asked to fixate at the centre of the screen, while the stimulation occurs on the edge of the visual field. These stimuli are composed of rotating or expanding wedges, made of checkerboard patterns. The study will last about one hour, depending on how long participants will choose to rest in between blocks. Participants will be given the opportunity to relax between blocks of presentation.

EEG and fMRI data will be recorded simultaneously, following guidelines and standard practices in place at the Centre for Integrative Neuroscience and Neurodynamics. Participants will be fitted with a MR-safe EEG system and comfortably positioned inside the MR scanner. Additional short and routine structural and functional scans will be performed before and after the blocks with visual stimulation, and will be used to assist in the statistical analyses of the data collected in this study.

2.3

Location

Post-recording interviews, instruction and debriefing will be performed in the psychology building on the University of Reading, Whiteknights campus. fMRI and EEG data collection will take place at the Centre for Integrative Neuroscience and Neurodynamics (CINN), also within the psychology building on the University of Reading, Whiteknights campus.

2.4

Funding

The research is currently funded by an EPSRC project EP/J003077/1, which covers costs of recording

2.5

Ethical Issues

There are no clinical physicians (i.e., neurologists, radiologists) taking part in this study. However, if the study investigators identify what they believe to be a significant abnormality, the study participant will be notified and may request that the MRI data be forwarded to his/her primary care physician. There will be no fees associated with this process.

EEG recording is generally considered risk and pain free, as is fMRI. However, if a participant feels uncomfortable and requests to cease study operations, we will do so immediately at no detriment to the participant. For instance, while claustrophobia is contraindicative of MRI study participation, if a participant felt claustrophobic while inside of the magnet, they will have easy access to an emergency alarm button, which will sound an alarm in the control room indicating that there is a problem. We will stop scanning immediately, and communicate with the patient to find out what the problem is. If they wish to be removed from the scanner, this can be done in a matter of seconds.

2.6 **Deception** The research project does not involve any deceptive strategy.
2.7 **Payment** Participants who are not taking part for Research Panel credits (students from Psychology) will be paid £10 per hour of study participation to the study. This is justified as the study will take over an hour per person, including instruction, setting up, experimental procedure and debriefing.
2.8 **Data storage, data protection and confidentiality** Each participant will be assigned a number, and all data will be referred to by the number and never associated with the actual name. Data will be stored on secure hard drives on University Computers, and will be password protected. Paper documents will be kept in a locked filing cabinet within the CINN. All information collected for the project will be kept for 5 years and then destroyed.
2.9 **Consent** Participants will be given the Information Sheet to read before taking part, and asked if they have any further questions or would like anything clarified. If they wish to proceed with the study, they will then be asked to sign the Consent Form (appended to this application).

Application UREC.pdf Application UREC.pdf Application UREC.pdf

SECTION 3: PARTICIPANT DETAILS
3.1 **Sample Size** We would like to record data from up to 30 participants. Previous visual stimulation studies using EEG or fMRI alone have shown that samples size of ten-fifteen and above are sufficient to obtain useful analysable data. However, combining the two modalities increases the number of artifacts in both data sets, meaning that most concurrent EEG-fMRI studies require around 20 participants. We therefore aim to record up to 30 participants in order to ensure that we have enough usable data sets.
3.2 Will the research involve vulnerable adults (e.g., adults with mental health problems or neurological conditions)? no
3.3 Will your research involve young people under the age of 18 years? Will your research involve children under the age of 5 years? **no**
3.4 Will your research involve NHS patients or Clients of Social Services? no
3.5 **Recruitment** Participants will be recruited from the student population at the University of Reading. Interested participants will be screened against any psychological or physical issue that may make their taking part in the study hazardous; e.g. epilepsy.

Application UREC.pdf Application UREC.pdf Application UREC.pdf

University of Reading

School of Systems Engineering

Researcher (principal): Dr. Etienne B. Roesch

Appendix A

INFORMATION SHEET #1
About the study

Background

The purpose of this study is the investigation of the neural circuitry involved in visual perception. In this study, we would like to investigate the neural activity evoked when one is presented with rotating and expanding wedges.

Should you choose to take part in this study, you will be fitted with our EEG system and installed in the MR scanner, and presented with visual stimuli composed of moving checkerboard patterns. Your task will simply be to fixate at the centre of the screen.

The experiment should not last more than 2 hours, with no longer than 1 hour inside the MR scanner. You will be given questionnaires to fill in before the experiment, and we will carry out a short debriefing interview after the experiment.

Potential discomforts

If you wish to drop out of the study at any time, you may do so with absolutely no penalty. If you feel uncomfortable, notify the researchers. If you wish to quit at any time, notify the researchers immediately.

Data confidentiality, storage and disposal of personal information

All data will be kept confidential and securely stored, with only a number attached to each participant, and therefore it will not be possible to link any set of data with any individual. In accordance with departmental policy, copies of your consent forms and screening forms and experimental data shall be kept for 5 years before they are destroyed.

If you wish to know the general results of the study please feel free to request this from the researcher who will provide you with a summary of the results when data analysis is complete.

Please Note

There is no intended clinical benefit to you from taking part in this study.

In exchange for your participation you shall receive payment of £10, unless you have been selected from the undergraduate research panel in which case you will receive a course credit. All participants can be emailed an image of their brain from the scan if they wish to receive one.

Taking part in this study is completely voluntary; you may withdraw at any time without having to give any reason. Should you withdraw from the study before its completion then you will still be entitled to payment/course credits.

This application has been reviewed by the University Research Ethics Committee and has been given a favourable ethical opinion for conduct.

Thank you for your help.

Application UREC.pdf Application UREC.pdf Application UREC.pdf

University of Reading

School of Systems Engineering

Researcher (principal): Dr. Etienne B. Roesch

INFORMATION SHEET #2
General Information on MRI

These notes give some information about an MRI study in which you are invited to take part.

Brief summary of the MRI technique

MRI is a method for producing images of the brain. It involves placing the participant inside a large, powerful magnet that forms part of the brain scanner. We use MRI to image the composition of different parts of your brain. We can also image which parts of your brain are more or less active. When particular regions of the brain are active, they require more oxygen, which comes from red corpuscles in the blood. As a result, the flow of blood increases. This can be detected as changes in the echoes from brief pulses of radio waves. These changes can then be converted by a computer into 3D images. This enables us to determine which parts of the brain are active during different tasks.

Potential risks with the procedure and precautions taken

As far as we know, this procedure poses no direct health risks. However, the Department of Health advises that certain people should NOT be scanned. Because the scanner magnet is very powerful, it can interfere with heart pacemakers and clips or other metal items which have been implanted into the body by a surgeon, or with body-piercing items. If you have had surgery which may have involved the use of metal items you should NOT take part. Note that only ferro-magnetic materials (e.g. steel) are likely to cause significant problems. Thus normal dental amalgam fillings do not prohibit you from being scanned, though a dental plate which contained metal would do so, and you would be asked to remove it. If you wear transdermal patches (e.g. nicotine patches) that contain metal, you will be asked to remove them. If you cannot, or are unwilling, to

remove these, please do not volunteer for this study. You will also be asked to remove metal from your pockets (coins, keys), remove articles of clothing which have metal fasteners (belts, bras, etc), as well as most jewellery. Alternative clothing will be provided as necessary. Watches and credit cards should not be taken into the scanner since it can interfere with their operation. We will provide you with a box in which you can leave your personal belongings securely while you are in the scanner. You will be

asked to complete a questionnaire (the Initial Screening Form) which asks about these and other matters to determine whether it is safe for you to be scanned. You will also be asked to complete a second, shorter, screening form immediately before the scan.

Possible discovery of findings related to research scans

There is no intended clinical benefit to you from taking part in this study. The scans are not intended to provide a medical diagnosis or a clean 'bill of health' – and the person conducting your scans will not be able to comment on the results of your scans.

The researchers involved do not have expertise in MRI diagnosis, as they are psychologists or allied scientists and are not medical doctors. We ask you to give the name and address of your general practitioner (GP). This is because occasionally, when we image healthy participants, the researchers may be concerned that a potential abnormality may exist on the scan. In such cases, we will send a copy of the image first to a radiologist for initial advice, then if advised to do so by the radiologist, to your GP. These qualified persons then decide what course of action is best. By signing the consent form, you authorise us to do this. If you are not willing to authorise this, please do not volunteer for the study.

It is important that you realise that these research scans are NOT a medical screening procedure, and will not provide any information that may help in the diagnosis of any medical condition. Even if the scan is reviewed by qualified experts, an additional medical issue cannot be ruled out. This is because the types of scans that we obtain for research purposes may be different to the scans used for diagnostic purposes. If you do have any health concerns, you should contact a qualified medical practitioner in the normal way.

What can you expect in the scanner

To be scanned, you would lie on your back on a narrow bed on runners, on which you would be moved until your head was inside the magnet. This is rather like having your head put inside the drum of a very large front-loading washing machine. The scanning process itself creates intermittent loud noises, and you would wear ear-plugs or sound-attenuating headphones. We would be able to talk to you while you are in the scanner through an intercom. If you are likely to become very uneasy in this relatively confined

part and this happens, you will be able to alert the researchers by activating an alarm and will then be removed from the scanner quickly. It is important that you keep your head as still as possible during the scan, and to help you with this, your head will be partially restrained with padded headrests. We shall ask you to relax your head and keep it still for a period that depends on the experiment but may be more than one hour (though you will not be continuously scanned for this period), which may require some effort on your part. If this becomes unacceptably difficult or uncomfortable, you may demand to be removed from the scanner.

You may be asked to look at a screen through a small mirror (or other optical device) placed just above your eyes. Details of the specific experiment in which you are invited to participate will either be appended to this sheet or else given to you verbally by the researcher. Detailed instructions will be given just before the scan, and from time to time during it.

Your participation is voluntary

You will be able to say that you wish to stop the testing and leave at any time, without giving a reason. This would not affect your relationship with the researchers in any way. The study will not benefit you directly, and does not form part of any medical diagnosis or treatment. If you agree to participate you will be asked to sign the initial screening form that accompanies this information sheet, in the presence of the researcher. It is perfectly in order for you to take time to consider whether to participate, or discuss the study with other people, before signing. After signing, you will still have the right to withdraw at any time before or during the experiment, without giving a reason.

How we handle confidential information

The images of your brain will be held securely and you will not be identified by name in any publications that might arise from the study. The information in the two screening forms will also be treated as strictly confidential and the forms will be held securely until eventually destroyed.

Further information about the specific study in which you are invited to participate may have been appended overleaf, if the researcher has felt that this would be helpful. Otherwise, he/she will already have told you about the study and will give full instructions prior to the scan. Please feel free to ask any questions about any aspect of the study or the scanning procedure before completing the initial screening form.

Application UREC.pdf Application UREC.pdf Application UREC.pdf

School University of Reading

Researcher (principal): Dr. Etienne B. Roesch

INFORMATION SHEET #3
General Information on EEG

What is an EEG?

EEG stands for electroencephalogram. It is a recording of the 'brainwaves' – the electrical activity of the brain. The many nerve cells that make up the brain produce continuous electrical activity when a person is awake, asleep or even in a coma. This can be recorded using small metal discs called electrodes, which are placed on the scalp. The electrical signals are then amplified by specialized equipment to produce what is seen in the EEG tracing. The signals show up on the EEG tracing as wavy lines, representing the fluctuations in electrical activity from moment to moment. Examining these tracings provide a lot of information about the workings of the brain.

What can I expect during an EEG?

After the EEG technician has asked you a few questions about your medical history, they will explain what will happen during the experiment. Once the EEG is fully set up, you will be placed inside the MRI scanner.

Preparation for an EEG

The first step is to apply the electrodes to the scalp. These are placed in standard positions, according to an internationally agreed convention. The technician will begin by measuring around the head with a tape measure to determine the optimal placement of the cap that hosts the electrodes. The skin where each electrode is to be placed will be cleaned to ensure that the electrical contact is good enough to allow the weak signals from the brain to be recorded properly. You will usually find that your appointment letter asks you to come for your EEG with 'clean, dry hair'. The request is made because

hair products and even the hairs natural oils can make it difficult for the technician to get a good electrode contact with the scalp. The electrodes, which are like little cup-shaped discs, are usually stuck in place with a special paste that also helps to conduct the electrical signals. Some of them may have a special conducting jelly put in them before they are stuck onto the scalp. The technicians are very skilled at doing this and it is not uncomfortable.

What happens during the EEG?

When everything is ready, the technician will ask you to lie in as relaxed a state as you can. This can be quite difficult if you're feeling nervous in unfamiliar surroundings, but once the test is underway, people usually find that they are gradually able to settle down. It is quite important to relax, since a tense person will have tense muscles around the face and forehead and these will produce electrical signals of their own. The electrodes on the scalp will pick these signals up too, obscuring the EEG and making the recording less useful.

What happens after the EEG?

The technician will carefully remove all the electrodes from your scalp. There are no after effects from the EEG so you can carry on with your normal activities immediately.

After the experiment, you will be provided a clean towel, shampoo and shower facilities to remove any electro-conductive gel from your hair. A hair-dryer is also available if required.

Potential risks with the procedure and precautions taken

EEG is risk and pain free, however it is advised that people with scalp conditions such as psoriasis do not take part, as a small amount of alcohol needs to be applied to the scalp, which would cause irritation if any such condition is present. We also advise that people with braided, dreadlocked or very long hair do not take part.

As you are participating in a simultaneous EEG-fMRI study, you should know that the special electrode caps that are used for that purpose are tested to be safe and non-reactive to the magnetic field. There is no added risk to you through the combination of these techniques beyond the risks listed for each technique separately. However, EEG recording within the scanner is much more sensitive for movement. So for the purposes of data quality, we would appreciate if you could minimize your movement during the study, especially during the task while the scanner is running.

Your data will be kept confidential and securely stored, with only a number attached to each participant, and therefore it will not be possible to link any set of data with any individual. All information collected for the project will be kept for 5 years and then

techniques beyond the risks listed for each technique separately. However, EEG recording within the scanner is much more sensitive for movement. So for the purposes of data quality, we would appreciate if you could minimize your movement during the study, especially during the task while the scanner is running.

Your data will be kept confidential and securely stored, with only a number attached to each participant, and therefore it will not be possible to link any set of data with any individual. All information collected for the project will be kept for 5 years and then destroyed. Taking part in this study is completely voluntary; you may withdraw at any time without having to give any reason. Please feel free to ask any questions that you may have about this study at any point. This application has been reviewed by the University Research Ethics Committee and has been given a favourable ethical opinion for conduct.

Application UREC.pdf Application UREC.pdf Application UREC.pdf

University of Reading

Appendix B: CONSENT FORM

Consent Form

1. I have read and had explained to me by Dr. Etienne B. Roesch/Catriona Scrivener

   the accompanying Information Sheet relating to the project on:

   ...Visual Perception.....

2. I have had explained to me the purposes of the project and what will be required of me, and any questions I have had have been answered to my satisfaction. I agree to the arrangements described in the Information Sheet in so far as they relate to my participation.

3. I understand that participation is entirely voluntary and that I have the right to withdraw from the project any time, and that this will be without detriment.

4. a) I authorise the Investigator to consult my General Practitioner, and provide their name and address details overleaf.

   b) I authorise my General Practitioner to disclose any information which may be relevant to my proposed participation in the project.

5. This application has been reviewed by the University Research Ethics Committee and has been given a favourable ethical opinion for conduct.

6. I have received a copy of this Consent Form and of the accompanying Information Sheet.

Name: .....

Date of birth [if GP details given]: .....

Signed: .....

Date: .....

Date of birth [if GP details given]: .....

Signed: .....

Date: .....

GP details

Name: .....

Address: .....

.....

.....

Appendix G

Copyright Permissions

Forms indicating permission to reuse content for Figures 1.3, 1.4, and 1.5 are appended here.

Figure G.1 Permission for Figure 1.3.

ELSEVIER LICENSE
TERMS AND CONDITIONS

Jul 24, 2020

This Agreement between Mr. Asad Malik ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	4875441171978
License date	Jul 24, 2020
Licensed Content Publisher	Elsevier
Licensed Content Publication	NeuroImage
Licensed Content Title	The spatiospectral characterization of brain networks: Fusing concurrent EEG spectra and fMRI maps
Licensed Content Author	David A. Bridwell,Lei Wu,Tom Eichele,Vince D. Calhoun
Licensed Content Date	Apr 1, 2013
Licensed Content Volume	69
Licensed Content Issue	n/a
Licensed Content Pages	11
Start Page	101
End Page	111
Type of Use	reuse in a thesis/dissertation

Figure G.2 Permission for Figure 1.4 (page 1 of 7). Only the first two pages have been included as the rest contain the standard terms and conditions. The complete document can be requested from the author.

Portion	figures/tables/illustrations
Number of figures/tables/illustrations	1
Format	both print and electronic
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title	Evaluation of ICA and Parallel ICA for Extracting Source Information from Simulated EEG-fMRI Signals
Institution name	University of Reading
Expected presentation date	Sep 2020
Portions	1g
	Mr. Asad Malik
Requestor Location	
Publisher Tax ID	GB 494 6272 12
Total	0.00 GBP

Terms and Conditions

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at http://myaccount.copyright.com).

GENERAL TERMS

Figure G.3 Permission for Figure 1.4 (page 2 of 7). Only the first two pages have been included as the rest contain the standard terms and conditions. The complete document can be requested from the author.

ELSEVIER LICENSE
TERMS AND CONDITIONS

Jul 24, 2020

This Agreement between Mr. Asad Malik ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	4875450477947
License date	Jul 24, 2020
Licensed Content Publisher	Elsevier
Licensed Content Publication	International Journal of Psychophysiology
Licensed Content Title	Unmixing concurrent EEG-fMRI with parallel independent component analysis
Licensed Content Author	Tom Eichele,Vince D. Calhoun,Matthias Moosmann,Karsten Specht,Marijtje L.A. Jongsma,Rodrigo Quian Quiroga,Helge Nordby,Kenneth Hugdahl
Licensed Content Date	Mar 1, 2008
Licensed Content Volume	67
Licensed Content Issue	3
Licensed Content Pages	13
Start Page	222
End Page	234
Type of Use	reuse in a thesis/dissertation

Figure G.4 Permission for Figure 1.5 (page 1 of 7). Only the first two pages have been included as the rest contain the standard terms and conditions. The complete document can be requested from the author.

Portion	figures/tables/illustrations
Number of figures/tables/illustrations	1
Format	both print and electronic
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title	Evaluation of ICA and Parallel ICA for Extracting Source Information from Simulated EEG-fMRI Signals
Institution name	University of Reading
Expected presentation date	Sep 2020
Portions	Fig 4.
	Mr. Asad Malik
Requestor Location	
Publisher Tax ID	GB 494 6272 12
Total	0.00 GBP
Terms and Conditions	

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at http://myaccount.copyright.com).

Figure G.5 Permission for Figure 1.5 (page 2 of 7). Only the first two pages have been included as the rest contain the standard terms and conditions. The complete document can be requested from the author.