

Paradigm Free Mapping:
Detection and Characterization of
Single Trial fMRI BOLD Responses
without Prior Stimulus Information

by

César Caballero Gaudes

Thesis submitted to
The University of Nottingham
for the degree of
Doctor of Philosophy

July 2010

Abstract

The increased contrast to noise ratio available at Ultrahigh (7T) Magnetic Resonance Imaging (MRI) allows mapping in space and time the brain's response to single trial events with functional MRI (fMRI) based on the Blood Oxygenation Level Dependent (BOLD) contrast. This thesis primarily concerns with the development of techniques to detect and characterize single trial event-related BOLD responses without prior paradigm information, Paradigm Free Mapping, and assess variations in BOLD sensitivity across brain regions at high field fMRI.

Based on a linear haemodynamic response model, Paradigm Free Mapping (PFM) techniques rely on the deconvolution of the neuronal-related signal driving the BOLD effect using regularized least squares estimators. The first approach, named PFM, builds on the ridge regression estimator and spatio-temporal t -statistics to detect statistically significant changes in the deconvolved fMRI signal. The second method, Sparse PFM, benefits from subset selection features of the LASSO and Dantzig Selector estimators that automatically detect the single trial BOLD responses by promoting a sparse deconvolution of the signal. The third technique, Multicomponent PFM, exploits further the benefits of sparse estimation to decompose the fMRI signal into a haemodynamical component and a baseline component using the morphological component analysis algorithm.

These techniques were evaluated in simulations and experimental fMRI datasets, and the results were compared with well-established fMRI analysis methods. In particular, the methods developed here enabled the detection of single trial BOLD responses to visually-cued and self-paced finger tapping responses without prior information of the events. The potential application of Sparse PFM to identify interictal discharges in idiopathic generalized epilepsy was also investigated. Furthermore, Multicomponent PFM allowed us to extract cardiac and respiratory fluctuations of the signal without the need of physiological monitoring.

To sum up, this work demonstrates the feasibility to do single trial fMRI analysis without prior stimulus or physiological information using PFM techniques.

Acknowledgements

This thesis has been a long, hard personal project. Nevertheless, this piece of work could never have been possible without the help and support of my supervisors, collaborators, friends and family.

First and foremost I would like to express my biggest gratitude to Prof. Penny Gowland, a supervisor I have been honoured to work with and who gave me excellent support and guidance to develop my ideas. I would also like to thank Prof. Ian Dryden for his advice and rigorous analysis. All my gratitude to Dr. Natalia Petridou, who taught me everything I know about fMRI analysis and always was there even in the long distance, and Dr. Sue Francis, who was always open to solve any question or doubt. My PhD was funded by the Marie Curie Fellowship under the European Commission's FP6 programme. A big thank you to Dr. Li Bai who gave me the opportunity to work at Nottingham and to the senior staff who put together this challenging project called CMIAG. I would also like to thank my examiners, Dr. Denis Schlupeck and Dr. Will Penny, for their useful comments on this thesis.

Throughout this thesis I was also fortunate to collaborate with Dr. Andrew Bagshaw, who noticed the potential of applying our work in epilepsy. I would also like to thank Dr. François Lazeyras for his patience, creating the perfect atmosphere to complete this thesis, and Prof. Dimitri Van de Ville who gave me short but to the point suggestions that inspired my work. It will be a great pleasure to continue my research with them at Geneva.

Next, I would like to thank my colleagues at the SPMMRC and CMIAG (I am afraid too many to write down all your names here), who made my PhD time an incredible experience, full of wonderful and unforgettable moments. My heart needs to say a big thanks to my good friends Aaron (you will never walk alone!!), Andrea, Olivier and Rosa for being always there. A Galia por aguantar mis quebraderos de cabeza.

Me gustaría expresar mi más profundo agradecimiento a mis padres, mi hermano y toda mi familia por la educación, el amor y el apoyo que siempre me han dado.

Finalmente, a Maria, sencillamente esta tesis no habría sido posible sin su amor, apoyo e inspiración. **This thesis is dedicated to her.**

César Caballero Gaudes

July 2010

Contents

1	Introduction	1
1.1	Overview of this thesis	2
2	Nuclear Magnetic Resonance and Magnetic Resonance Imaging	3
2.1	Principles of Nuclear Magnetic Resonance	3
2.2	Principles of Magnetic Resonance Imaging	11
2.3	Conclusions	25
3	Functional Magnetic Resonance Imaging	27
3.1	Functional Imaging	27
3.2	BOLD Functional MRI	29
3.3	BOLD sensitivity in Gradient Echo Echo Planar Imaging	33
3.4	Noise in BOLD fMRI	35
3.5	FMRI data analysis	37
3.6	Correcting BOLD sensitivity maps	46
3.7	Conclusion	53
4	Paradigm Free Mapping	55
4.1	Introduction to Paradigm Free Mapping	55
4.2	Theory	59
4.3	Methods	70
4.4	Results	77
4.5	Discussion	90
4.6	Conclusion	98
5	Sparse Paradigm Free Mapping	99
5.1	Theory	99
5.2	Methods	107

5.3	Results	114
5.4	Discussion	132
5.5	Conclusion	138
6	Multicomponent Paradigm Free Mapping	139
6.1	Theory	139
6.2	Methods	146
6.3	Results	149
6.4	Discussion	157
6.5	Conclusion	164
7	Conclusions and further work	165
7.1	Summary and final discussion	165
7.2	Further developments	168
7.3	List of publications	170
	Appendices	173
A.1	Wiener Deconvolution	173
A.2	Prewhitening	174
A.3	Levinson Durbin algorithm	174
A.4	Autocorrelation coefficients from the AR and reflection coefficients .	177
A.5	Number of degrees of freedom of the t -statistic	177
	Bibliography	181

List of Acronyms

AIC	Akaike Information Criterion
AICc	Finite Sample AIC
AR	Autoregressive
ARMA	Autoregressive Moving Average
ASL	Arterial Spin Labeling
ATS	Activation Time Series
BIC	Bayesian Information Criterion
BH	Benjamini-Hochberg FDR procedure
BOLD	Blood Oxygenation Level Dependent
CBF	Cerebral Blood Flow
CBV	Cerebral Blood Volume
CMRO₂	Cerebral Metabolic Rate of Oxygen
CNR	Contrast to Noise Ratio
CSF	Cerebrospinal Fluid
DASSO	Dantzig Selector with Sequential Optimization
DS	Dantzig Selector
EEG	Electroencephalography
EPI	Echo Planar Imaging
EM	Electromagnetic
FDR	False Discovery Rate
FID	Free Induction Decay
FIR	Finite Impulse Response
FLASH	Fast Low Angle Shot

fMRI	Functional Magnetic Resonance Imaging
FN	False Negative
FNR	False Negative Rate
FP	False Positive
FPR	False Positive Rate
GE	Gradient Echo
GLM	General Linear Model
GLS	Generalized Least Squares
GM	Gray Matter
HRF	Haemodynamic Response Function
IED	Interictal Epileptic Discharge
IR	Inversion Recovery
LARS	Least Angle Regression
LASSO	Least Absolute Shrinkage and Selection Operator
LD	Levinson Durbin Recursion
LS	Least Squares
LUT	Lower Universal Threshold
MA	Moving Average
MAP	Maximum a-posteriori
MCA	Morphological Component Analysis
MDL	Minimum Description Length
MDL_c	Finite Sample MDL
MEG	Magnetoencephalography
MPRAGE	Magnetization Prepared Rapid Gradient Echo
MRI	Magnetic Resonance Imaging
NMR	Nuclear Magnetic Resonance
P_{ET}CO₂	Pressure of End Tidal CO ₂
PFM	Paradigm Free Mapping
RETROICOR	Retrospective Correction of Physiological Motion Effects in fMRI
RF	Radiofrequency

RR	Ridge Regression
SE	Spin Echo
SNR	Signal to Noise Ratio
SNR₀	Signal to Thermal Noise Ratio
SPFM	Sparse Paradigm Free Mapping
SPM	Statistical Parametric Map
SPT	Self Paced Tapping
SURE	Stein's Unbiased Risk Estimation
TE	Echo Time
TI	Inversion Time
TP	True Positives
TN	True Negatives
TR	Repetition Time or Temporal Resolution
tSNR	Temporal Signal to Noise Ratio
UT	Universal Threshold
VCT	Visual Cued Tapping
WM	White Matter

Chapter 1

Introduction

Since its introduction in the early 1990s, functional Magnetic Resonance Imaging (fMRI) has become the most widespread tool for neuroscientists, clinicians and physicists to investigate human brain function. fMRI is a noninvasive imaging technique that measures local changes in blood magnetic susceptibility related to changes in blood flow and oxygenation that occur due to increased neuronal activity. In particular, characterizing the haemodynamic response to short event-related stimuli enables the study of the temporal orchestration of cortical events with fMRI, also known as mental chronometry. Event-related fMRI analysis is usually performed by averaging the single trial response to multiple successive stimuli in order to improve the estimate of the haemodynamic parameters. With the advent of ultra high MR scanners, the contrast to noise ratio of the BOLD response is enhanced so that measuring the response to individual events becomes feasible. Single trial fMRI has the potential to study finer cognitive processes, such as learning or adaptation, or the results of pharmacological perturbations. Assuming *a priori* knowledge of the timing of the events is nowadays the prevalent approach for single trial fMRI analysis. However, this information can be difficult to obtain accurately under certain clinical, behavioural or experimental circumstances. For instance, recording the timing of perceptual changes in binocular rivalry requires subject's feedback, or the onset of interictal epileptic discharges is obtained from simultaneous EEG data in EEG-fMRI studies in epilepsy. Therefore, there is a need for techniques to fMRI data analysis which avoid the requirement to specify the onsets of cortical responses, and which allow the study and use of more unconstrained experimental paradigms, particularly for single-trial events. This thesis aims to develop novel analysis tools to enable the detection of single trial fMRI BOLD responses without prior stimulus information.

1.1 Overview of this thesis

The second chapter of this thesis describes the fundamental principles of nuclear magnetic resonance and magnetic resonance imaging. An overview of the main MR pulse sequences used in this thesis is also given.

Within the framework of functional brain imaging, Chapter 3 describes the underpinnings of fMRI based on the blood oxygenation level dependent (BOLD) contrast. The description focus on the analysis of BOLD sensitivity in gradient echo fMRI in order to describe how to generate sensitivity maps of BOLD detection and formulates a statistical correction procedure to compensate for sensitivity variations across brain regions.

Chapters 4, 5 and 6 describe three novel methods to detect and characterize single trial BOLD responses without prior information of the event timing: Paradigm Free Mapping.

Chapter 4 begins with a description of the rationale of Paradigm Free Mapping (PFM), and subsequently introduces the first PFM methodology based on the ridge regression deconvolution of the BOLD haemodynamic response and statistical inference based on spatio-temporal t -statistics against a baseline period. The method is tested on a visuomotor paradigm for paradigm free detection of single trial BOLD responses.

Chapter 5 investigates the use of two sparse estimators, the Least Absolute Shrinkage and Selection Operator (LASSO) and the Dantzig Selector, for paradigm free mapping. The potential application of Sparse Paradigm Free Mapping to detect interictal epileptiform discharges without requiring EEG-based information is also studied.

Chapter 6 extends the principles of sparse estimation in order to formulate a Multicomponent PFM approach which simultaneously decomposes fMRI voxel time series into two components, a haemodynamic component with single trial BOLD responses and a component modelling baseline fluctuations with a sinusoidal basis set. The usefulness of this approach is illustrated by comparing it with the Retrospective Correction of Physiological Motion Effects (RETROICOR) method since the technique enables the correction of cardiac and respiratory fluctuations without the need of physiological monitoring.

Finally, Chapter 7 summarizes the findings arising from the work developed in this thesis and outlines research topics for further development.

Chapter 2

Nuclear Magnetic Resonance and Magnetic Resonance Imaging

This chapter provides an introduction to the basic features of Nuclear Magnetic Resonance (NMR) and Magnetic Resonance Imaging (MRI). Understanding the basic principles of both techniques is pertinent in order to further comprehend the physical underpinnings of the signal observed in functional MRI experiments. Firstly, in §2.1, the principles of NMR are briefly described. This is followed in §2.2 by a description of how the NMR signal can be exploited for the imaging of objects with MRI.

The basic principles of NMR and MRI are described in multiple excellent textbooks. The review presented here is mainly based on [34, 59, 184, 274].

2.1 Principles of Nuclear Magnetic Resonance

Nuclear Magnetic Resonance (NMR) is the study of the dynamics of magnetic nuclei in interaction with an external magnetic field [294]. NMR is the basic physical principle behind Magnetic Resonance Imaging and its basic theory is described in this section.

2.1.1 Interaction of Nuclear Spins with Magnetic Fields

Subatomic particles in an atomic nucleus, i.e. protons and neutrons, have an intrinsic angular momentum due to the rotation along its axis, which is usually referred to as *spin*. Contrary to classical spinning systems, spins of protons and neutrons can only have an angular momentum with discrete magnitude, the spin quantum number \mathbf{I} , defined by quantum mechanics [184].

In order to form an atomic nucleus, protons and neutrons combine with other equivalent spins of opposite states which causes the net spin magnitude to be zero for atoms with an even number of protons or neutrons, or different from zero for atoms with an odd number of protons or neutrons [59]. For example, the most widespread nuclei studied in NMR is hydrogen, ^1H , which consists of a single proton and has a net spin of $\mathbf{I} = \frac{1}{2}$ producing two possible alignments $\mathbf{M} = \pm\frac{1}{2}$.

As the atomic nucleus has an electric charge, the rotating spin generates a small but noticeable magnetic dipole moment. When the spin is located within an external magnetic field, B_0 , it will align either in the same or opposite direction of the field depending on the state of the spin. In the absence of thermal agitation, all spins would align in the direction of the magnetic field. Nevertheless, when a sample is at physiological temperatures, not all the spins align with the magnetic field, but the number of spins aligning in the direction of B_0 is slightly larger than those aligning in the opposite direction (Figure 2.1).

Strictly speaking, since the spin has an angular momentum, its spin axis does not exactly align with the field but *precesses* around the field axis at a specific frequency, called the Larmor frequency,

$$\omega_0 = \gamma B_0, \quad (2.1)$$

where γ is the gyromagnetic ratio in radians/Tesla and B_0 is the magnitude of magnetic field in Tesla¹. The value of γ depends on the nature of the nuclei, e.g. $\gamma = 2.675 \times 10^8 \text{ s}^{-1}\text{T}^{-1}$ (or $\frac{\gamma}{2\pi} = 42.57 \text{ MHz/T}$) for ^1H . If electro-magnetic (EM) radio-frequency (RF) pulses are applied with frequency equal to the Larmor frequency, i.e. at *resonance*, a transition between the energy states of spins is produced [42, 293].

In practice, in order to simplify the analysis of the NMR experiment, the combined effect of the whole ensemble of spins is generally considered rather than study-

¹The strength of Earth's magnetic field at 0° is $31 \mu\text{T}$ ($3.1 \times 10^{-5}\text{T}$)

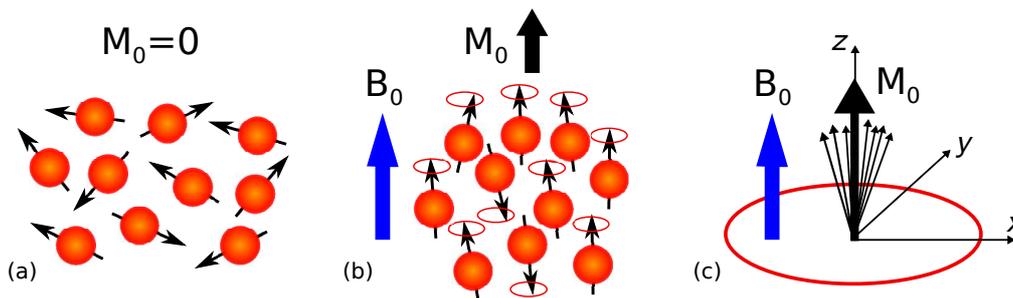


Figure 2.1: Spin interaction with a strong magnetic field. (a) The nuclear magnetic moments are initially randomly oriented. (b) When located in a strong magnetic field B_0 , spins align parallel or antiparallel to the direction of field and precess at the Larmor frequency, (c) creating a bulk magnetization moment M_0 in the direction of B_0 . Adapted from [184].

ing the effect of the magnetic field in each particular spin. As a whole, the excess in the number of protons which align with B_0 yields a bulk magnetization moment, M_0 , different from zero in the direction of field, conventionally the z -direction (Figure 2.1c). This induced magnetization due to the static magnetic field is often referred to as equilibrium, net or bulk magnetization, and it is proportional to the excess difference and B_0 .

2.1.2 The Effect of Radio-Frequency Pulses

Equilibrium magnetization produces no signal in the MR receivers since these only detect signal in the plane transverse to B_0 . If, in presence of the static magnetic field, an RF pulse B_1 is applied orthogonally to B_0 , i.e. in the x - or y -direction, with frequency equal to the Larmor frequency ω_0 , resonance occurs, the individual spins flip, but the net magnetization tips down an angle α towards the x - y plane or transverse plane that depends on the duration τ and amplitude of the B_1 RF pulse:

$$\alpha = \gamma B_1 \tau. \quad (2.2)$$

This process of flipping can be observed from two different reference frames, the laboratory frame and the rotating frame. When observing from the laboratory frame, one is observing the process from outside and therefore the magnetization would be precessing about the z -direction of B_0 whilst being tipped onto the x - y plane by B_1 . However, if one jumps on board of the magnetization vector M_0 , one is observing from the rotating frame and the magnetization vector only tips down in the z direction. The difference between both frames is illustrated in Figure 2.2.

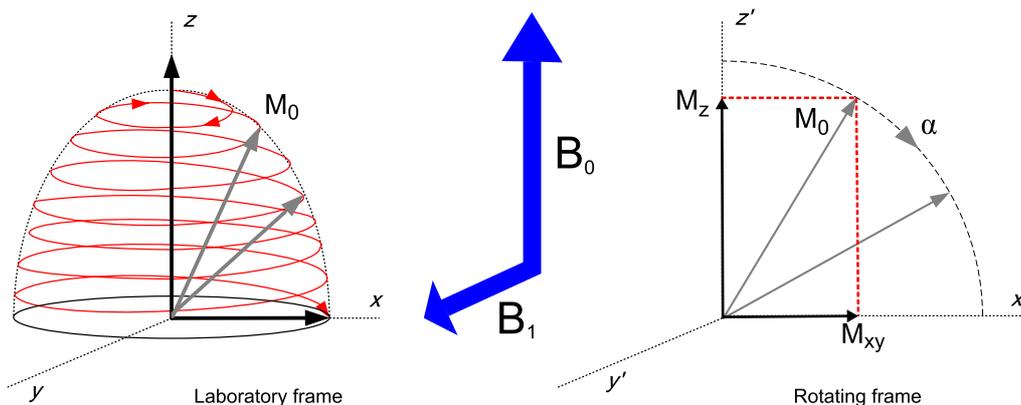


Figure 2.2: Evolution of the bulk magnetization vector M_0 when applying an RF pulse at the Larmor frequency, in the laboratory (x, y, z) and rotating frames (x', y', z') .

The magnetization vector in the transverse plane and the z direction are denoted as M_{xy} (transverse magnetization) and M_z (longitudinal magnetization), respectively. An α -degree pulse results in the net longitudinal, M_z , and transverse, M_{xy} , magnetization being

$$M_z(0) = M_0 \cos(\alpha) \quad (2.3)$$

$$M_{xy}(0) = M_0 \sin(\alpha) \quad (2.4)$$

For instance, if a 90° pulse is applied, the bulk magnetization M_0 flips into the transverse plane and the magnitude of M_{xy} is M_0 and M_z is zero. In this case, the magnetization is said to be excited or saturated and the 90° pulse is referred to as an excitation pulse. On the other hand, a 180° pulse will flip the magnetization to $-z$ direction, which is called inverted magnetization, and the 180° pulse is referred to as an inversion pulse.

2.1.3 Transverse and Longitudinal relaxations

Observing the magnetization from the laboratory frame, immediately after the application of an RF pulse the system will be at non-equilibrium and the spins will be rotating in phase around the z -axis and at the Larmor frequency in the transverse plane (Figure 2.3). But, just after the RF pulse is turned off, spins will begin returning to the equilibrium magnetization by spin-lattice relaxation processes and spins will begin dephasing because they precess with different frequencies due to spin-spin relaxation processes.

Spin-Lattice relaxation

Spin-Lattice relaxation or longitudinal relaxation is the process whereby energy absorbed by the spins is released back into the surrounding lattice, reestablishing the spin equilibrium distribution. Physically the main processes responsible for this phenomenon are dipole-dipole interactions. Motion of molecules can give rise to oscillating magnetic field in the environment of the spin under consideration. The dynamics of the magnetization are described by the Bloch equation [42]

$$\frac{d\mathbf{M}}{dt} = \gamma(\mathbf{M} \times \mathbf{B}) - \frac{(M_z - M_0)\mathbf{k}}{T_1} - \frac{M_x\mathbf{i} + M_y\mathbf{j}}{T_2}. \quad (2.5)$$

Using (2.3), the evolution of the longitudinal magnetization M_z is given by

$$M_z(t) = M_0 \left[1 - (1 - \cos(\alpha)) e^{-t/T_1} \right], \quad (2.6)$$

where t is the time after the RF pulse has been applied. The longitudinal magnetization M_z exponentially recovers into $M_z(0)$ along the z -axis with a spin-lattice relaxation time constant T_1 (Figure 2.4a). The value of T_1 depends on the tissue composition, temperature and magnetic field strength [375].

Spin-Spin relaxation

After the excitation RF pulse, all spins are in phase and precessing at the same frequency. All the spins precessing in phase give rise to a rotating net transverse magnetization. T_1 recovery will lead to a reduction in this signal but it is also reduced by any process that leads to spin dephasing. Again magnetic field from neighbouring molecules can interact with applied magnetic field. If the magnetic field generated by a neighbouring proton increases (or decreases) the applied field, then the field experienced by the spin also increases (or decreases). This spin-spin interaction forces the Larmor frequency to vary across the sample which makes every spin precess at a different frequency in the rotating frame (Figure 2.3).

Spin dephasing causes an exponential decay in the transverse magnetization M_{xy} with a spin-spin relaxation time constant T_2 , also called transverse relaxation time constant (Figure 2.4b). The dynamics of the transverse magnetization are also

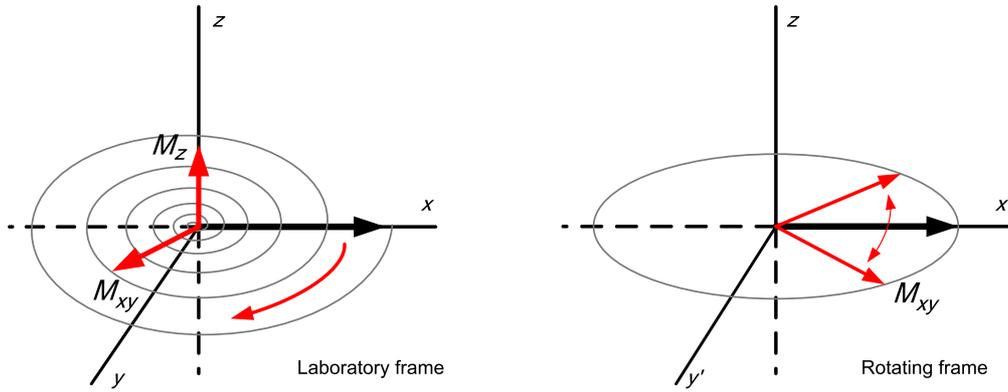


Figure 2.3: Decay of the transverse magnetization due to spin-spin interactions.

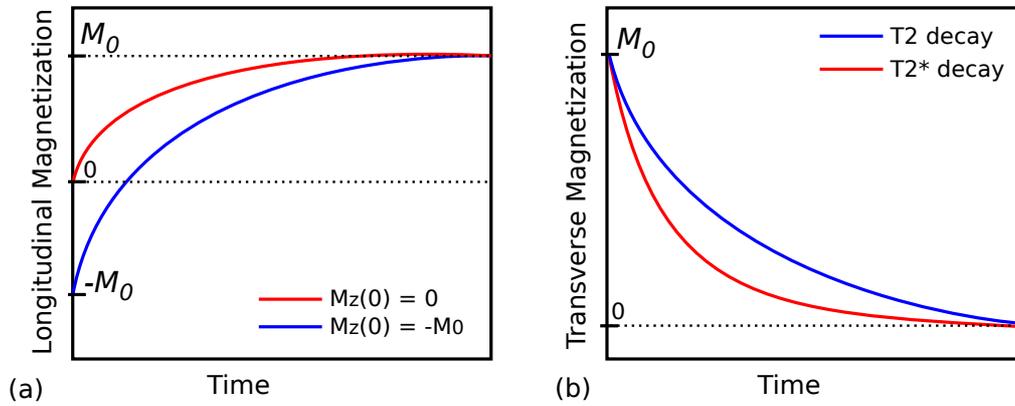


Figure 2.4: (a) Spin-Lattice relaxation of the longitudinal magnetization, M_z , with relaxation time constant T_1 . When the flip angle $\alpha = 90^\circ$, $M_z(0) = 0$ (red line) and when $\alpha = 180^\circ$, $M_z(0) = -M_0$ (blue line). (b) Spin-Spin Relaxation of the transverse magnetization, M_{xy} . The T_2^* -decay (red line) is faster than the T_2 -decay (blue line). Adapted from [274].

defined from the Bloch equation (2.5) given by

$$M_{xy}(t) = M_0 \sin(\alpha) e^{-t/T_2}. \quad (2.7)$$

In practice, the decay of M_{xy} is faster with a true relaxation time constant T_2^* (Figure 2.4b). This is a combination of true T_2 spin-spin interactions that are stochastic and therefore irreversible, and T_2' decay due to B_0 inhomogeneities which can be measured and reversed. In sum, the transverse relaxation time is given by

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}. \quad (2.8)$$

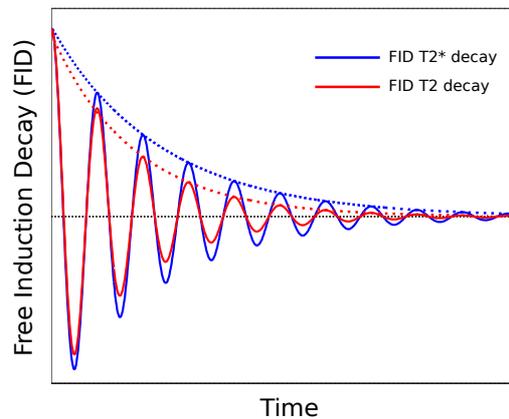


Figure 2.5: Free induction decay signal for the T_2 and T_2^* decays.

2.1.4 The Free Induction Decay Signal

The energy absorbed by the spins due to the RF pulse is retransmitted by means of EM radiation at frequencies close to the Larmor frequency. The MR spectrometer captures the energy related to the transverse magnetization, M_{xy} , by means of a phase sensitive detector, which then decouples the received field in-phase and quadrature components and demodulates the signal in frequency by multiplying with a reference signal with the same frequency as the transmitted RF pulse. After low-pass filtering, the complex baseband signal has a bandwidth equal to the range of Larmor frequencies across the sample. This signal is called the Free Induction Decay (FID) signal (Figure 2.5) and its amplitude depends on the longitudinal and transverse relaxation processes described above.

2.1.5 Echoes

The amplitude of the FID signal decays very rapidly with time constant T_2^* . Therefore, some kind of trick is necessary to revert the dephasing process and increase again the amplitude of the observed signal. Let us assume that a spin precesses at a constant frequency and accumulates a phase ϕ after a period of time $t = TE/2$. Within a sample, some spins will precess at faster or slower frequencies than others and therefore the 'faster' spins would have accumulated a higher phase during $TE/2$. If a 180° RF pulse is applied at that particular time $TE/2$, the spin phases will be flipped to $-\phi$. As the spin precesses at the same frequency it will again accumulate a phase ϕ for $TE/2$ but in opposite direction. Faster spins will accumulate more

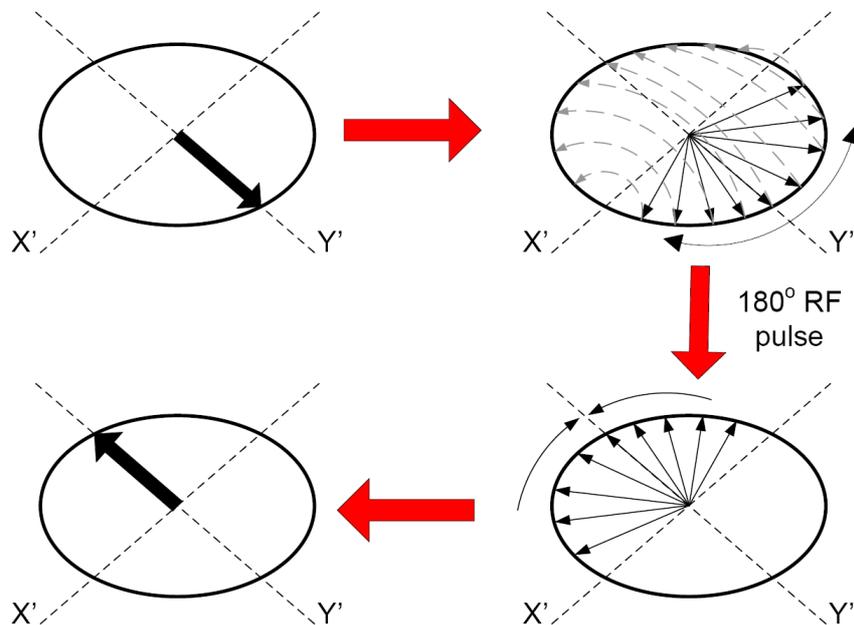


Figure 2.6: Formation of an echo. After tipping the magnetization into the transverse plane, spins precess at different frequencies and the magnetization vectors begin to dephase. The 180° pulse at time $TE/2$ flips the transverse plane into the other direction, and each magnetization vector continues precessing at the same frequency so that they refocus to form an echo at the echo time TE .

phase causing net phase coherence at $t = TE$, i.e. the phase of all spins will be ideally zero. Because of this effect the 180° pulse is often called a refocusing pulse and the process is known as phase refocusing. This recovery in phase coherence of the spins causes an increase in the FID signal, which is known as an echo, at time TE , which is known as the Echo Time [157]. The process of echo generation is shown in Figure 2.6.

As described in §2.2.4, when an echo is created by means of a refocusing pulse it is called spin echo, whereas if it is created by means of magnetic field gradients, it is called gradient echo. As shown in Figure 2.7, this process can be repeated multiple times such that successive refocusing pulses produces multiple echoes. The only limitation is the decrease in the signal amplitude due to the natural T_2 decay which always exists for both spin and gradient echo techniques. Nevertheless, contrary to gradient echoes, spin echoes are insensitive to B_0 field inhomogeneity leading to signal losses in inhomogeneous areas near air-tissue or bone-tissue boundaries.

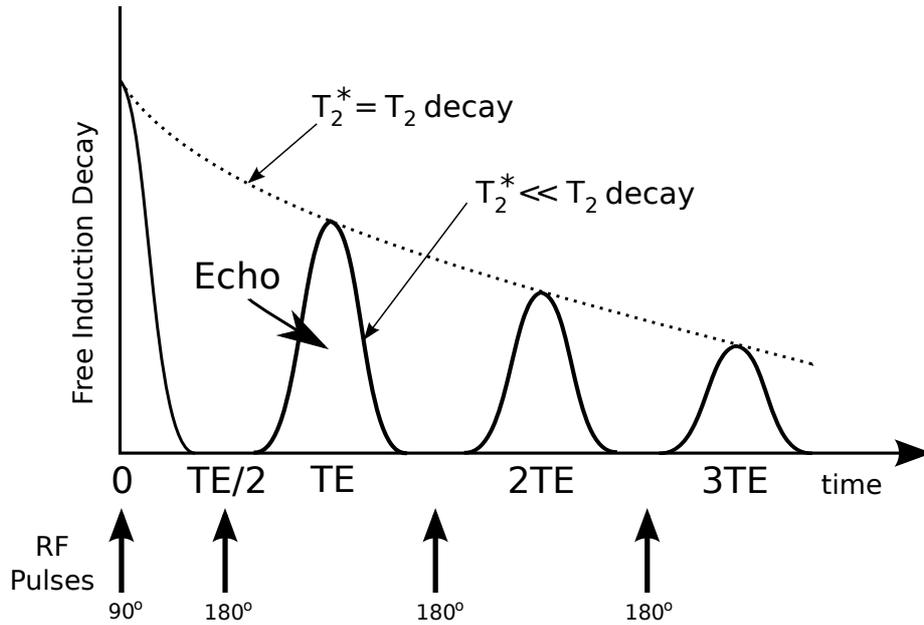


Figure 2.7: Generation of multiple spin echoes with RF refocusing pulses. Adapted from [59].

2.2 Principles of Magnetic Resonance Imaging

Prior the 1970s NMR was mainly used to study the chemical properties of samples with magnetic resonance spectroscopy [112]. In 1973, Lauterbur [205] and Mansfield [237] proposed that NMR principles could also be employed to obtain 2D MR images of a sample by spatially varying the applied magnetic field across the sample. Those first experiments gave birth to Magnetic Resonance Imaging [106, 169, 202, 235, 238], whose principles are described in this section.

2.2.1 Signal Localization with Magnetic Field Gradients

The NMR signal must be acquired at different spatial locations to generate an MR image. Modern MRI scanners use variations in the applied magnetic field so that the observed Larmor frequency is spatially dependent. This is done by modifying the frequency of precession at different locations with the addition of small spatially-controlled gradients of the static B_0 magnetic field. Generally, linear field gradients are used to encode the spatial location in each direction,

$$G_x = \frac{dB_x}{dx}, \quad G_y = \frac{dB_y}{dy}, \quad G_z = \frac{dB_z}{dz}. \quad (2.9)$$

Hence, in the presence of a magnetic field gradient the local applied magnetic field at point $\mathbf{r} = (x, y, z)$ is given by

$$\begin{aligned} B(\mathbf{r}) &= B_0 + xG_x\mathbf{i} + yG_y\mathbf{j} + zG_z\mathbf{k} \\ &= B_0 + \mathbf{G}_r \cdot \mathbf{r}, \end{aligned} \quad (2.10)$$

where $\mathbf{G}_r = (G_x, G_y, G_z)$ is the applied gradient in mT/m units from the gradient RF coils. The absolute amplitude of the field increases proportionally to the distance of r from a reference point, conventionally set at the center of the magnet. According to the Larmor equation (2.1), the resonance frequency of each spin after an excitation pulse would also depend on its position \mathbf{r} in the sample as follows,

$$\begin{aligned} \omega(\mathbf{r}) &= \gamma B(\mathbf{r}) \\ &= \gamma (B_0 + \mathbf{G}_r \cdot \mathbf{r}). \end{aligned} \quad (2.11)$$

As a result, when an RF pulse with a bandwidth $\Delta\omega$ is applied in presence of a magnetic field gradient \mathbf{G}_r , the sample is excited within an specific region. Therefore, differences in the resonance frequency due to magnetic field gradients can be scattered across the bandwidth of the FID signal and this frequency diversity can be used to encode different spatial locations.

The following sections describe three methods for signal localization. Section §2.2.2 introduces the slice selection concept for selective excitation and §2.2.3 introduce the spatial localization of the signal in a 2D Fourier space, usually known as k -space.

2.2.2 Slice selection

Instead of exciting the whole sample with an RF non-selective pulse, as described so far, it is usual in MRI to selectively excite the spins within a cross-sectional slice of the sample [139]. Selective excitation is carried out by applying an RF pulse with a linear magnetic field gradient in the direction of the slices to be excited. This gradient field sets the Larmor frequency of the spins within the slice to match the frequencies centered in the RF pulse.

For instance, let us assume that slice selection is done in the z -direction (Figure

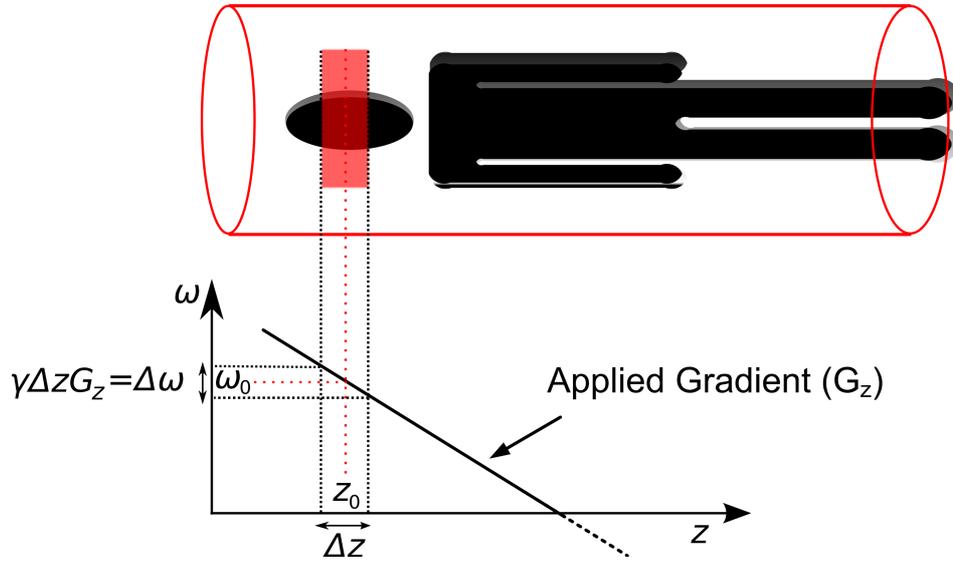


Figure 2.8: Slice selection with magnetic field gradients. When an RF pulse with a range of frequencies, $\Delta\omega$, is applied in the presence of a magnetic field gradient, G_z , a slice of thickness, Δz is excited.

2.8). To excite a slice of width Δz a gradient field G_z must be applied to cause a linear variation of the Larmor frequency across the slice along with an RF pulse of carrier frequency ω_0 and bandwidth $\Delta\omega$. The slice thickness is then given by

$$\Delta z = \frac{\Delta\omega}{\gamma G_z}. \quad (2.12)$$

This pulse will only excite those spins within $z_0 + \frac{\Delta z}{2} > z > z_0 - \frac{\Delta z}{2}$, where the location of z_0 is defined by the carrier frequency of the pulse ω_0 with respect to the reference point of the magnet. Spins on either side of the slice will precess at different frequencies and are thus largely unaffected by the pulse in terms of being tipped. In order to obtain MR images with high spatial resolution in the slice direction we could increase the steepness of the gradient, which implies higher current and stronger fields, and/or decrease the bandwidth of the RF pulse, which implies longer pulse durations and therefore longer scanning times. In practice, some cross-talk into neighbouring slices always exists since an ideal square slice profile is not possible due to finite time duration of the pulse. The effect of crosstalk is usually overcome by acquiring interleaved slices or leaving a gap between slices.

2.2.3 Fourier Imaging: Phase Encoding, Frequency Encoding and the k -space

In the rotating reference frame, after applying an RF pulse the signal emanating from point \mathbf{r} in an object at time t is given by [235, 237]

$$S(\mathbf{r}, t) \propto \rho(\mathbf{r}) e^{-i \int_0^t \omega(\mathbf{r}, t') dt'}, \quad (2.13)$$

where $\rho(\mathbf{r})$ is the spin density distribution and the frequency of precession of the spins is $\omega(\mathbf{r}, t') = \omega_0 + \gamma \mathbf{r} \cdot \mathbf{G}(t')$, is conventionally defined in radians. Comparing with equation (2.11), the field gradients are not only spatial-dependent on \mathbf{r} but also time-dependent and so the signal depends on the evolution of the gradients over time. In practice, the FID signal is collected across an extended region of the sample,

$$S(t) \propto \int_{z_0 - \frac{\Delta z}{2}}^{z_0 + \frac{\Delta z}{2}} \int_{y_0 - \frac{\Delta y}{2}}^{y_0 + \frac{\Delta y}{2}} \int_{x_0 - \frac{\Delta x}{2}}^{x_0 + \frac{\Delta x}{2}} \rho(x, y, z) e^{-i \gamma \int_0^t (B_0 + \mathbf{r} \cdot \mathbf{G}(t')) dt'} dx dy dz, \quad (2.14)$$

where (x_0, y_0, z_0) and $(\Delta x, \Delta y, \Delta z)$ are the coordinates of the centre and dimensions of the field of view in each direction, respectively. First, a slice in the z -direction is excited during slice selection, as explained in §2.2.2. Next, gradient pulses are applied in the x - and y -direction in order to obtain a cross-sectional 2D image of the spin density distribution in the $x - y$ plane. After excitation the transverse magnetization will start to accrue a phase due to the gradients applied which at time t will be $\phi = \omega(x, y)t$ where

$$\omega(x, y) = \gamma(B_0 + G_x x + G_y y). \quad (2.15)$$

After demodulation to baseband by the Larmor frequency ($\omega_0 = \gamma B_0$), the FID signal from an (x, y) region is given by

$$S(t) \propto \int_{y_0 - \frac{\Delta y}{2}}^{y_0 + \frac{\Delta y}{2}} \int_{x_0 - \frac{\Delta x}{2}}^{x_0 + \frac{\Delta x}{2}} \rho(x, y) e^{-i \gamma (G_x x + G_y y) t} dx dy. \quad (2.16)$$

Defining the k -terms

$$\begin{aligned} k_x &= -\gamma G_x t \\ k_y &= -\gamma G_y t, \end{aligned} \quad (2.17)$$

the FID signal can be written as

$$S(k_x, k_y) \propto \int_y \int_x \rho(x, y) e^{i(k_x x + k_y y)} dx dy = \int_y \int_x \mathcal{F}^{-1} \{ \rho(x, y) \} dx dy. \quad (2.18)$$

This equation can be interpreted as the inverse Fourier transform of the spin density distribution $\rho(x, y)$ of the sample at position (x, y) , and the k_x and k_y terms can be thought as spatial frequencies. Consequently, an image of the spin density at position (x, y) can be obtained by applying the Fourier transform on the received FID signal as

$$\rho(x, y) \propto \int_{k_x} \int_{k_y} S(k_x, k_y) e^{-i(k_x x + k_y y)} dk_x dk_y = \int_{k_y} \int_{k_x} \mathcal{F} \{ S(k_x, k_y) \} dk_x dk_y. \quad (2.19)$$

This concept of Fourier imaging forms the basis of modern imaging techniques with MRI. The sampling of the spatial frequencies k_x and k_y is carried out with two encoding methods in the x -axis and y -axis known as *frequency encoding* and *phase encoding*, respectively.

Frequency and phase encoding

Frequency encoding and phase encoding are applied after slice selection during the image readout period in order to encode the 2D image in the frequency domain for those spins that have already been excited by slice selection.

Conventionally, frequency encoding occurs in the x -direction with a linear gradient G_x . Following (2.14) and (2.19), the FID signal due to this gradient will be

$$S(k_x) \propto \int_y \rho(x, y) e^{ik_x x} dy, \quad (2.20)$$

providing information about the spin density population in the x -direction. At this point, all the signals with a given x -coordinate contribute to the net signal. Between each line encoded in the x -direction, a linear magnetic field gradient G_y is applied

for a short interval to produce a small phase increment in the net signal and encode the image in the y -coordinate. This is called phase encoding [106, 202].

K-space

The concept of k -space was formulated in 1983 by Ljunggren [221] and Twieg [347] based on the ideas of Fourier imaging [237]. The spatial frequencies k_x and k_y facilitate the visualization of the dynamics of the spin phase under the influence of magnetic field gradients in the frequency space, commonly referred to as the k -space. In general, the k -space is defined by the spatial frequencies (k_x, k_y, k_z) corresponding to the spatial coordinates (x, y, z) . Normally, most MRI techniques are 2D planar imaging where only one slice is acquired at a time by keeping k_z fixed, i.e. selectively exciting just that one slice, and varying k_x and k_y (see Figure 2.9).

In 2D planar imaging, immediately after excitation of the slice with the RF pulse, the sampling position is at the centre of k -space and the sampling position at time t depends on the accumulated-phase of the magnetization due to the gradients along time. Low spatial frequencies are mapped in the central region of the k -space, and the edges correspond to high spatial frequencies, as shown in Figure 2.9. In order to fully sample the k -space the k -cursor is moved according to a k -space trajectory defined with the frequency encoding and phase encoding gradients. During frequency encoding, data points are recorded at regular intervals along a line in the k_x direction, whereas during phase encoding the spatial frequency k_y is altered without recording the signal.

The k -space represents a continuous range of frequencies, but the FID or echo signal is only sampled at discrete intervals. The spacing and location of the k -space measurements determine the field-of-view (FOV) and spatial resolution of the resulting image. The FOV is inversely proportional to the spacing of k -space samples which, in turn, depends on the difference in phase between two samples ($G\Delta t$). For example, in the frequency encoding direction,

$$\text{FOV}_x = \frac{1}{\Delta k_x} = \frac{1}{G_x \Delta t}, \quad (2.21)$$

and the largest value of k , k_{max} , determines the spatial resolution of the recon-

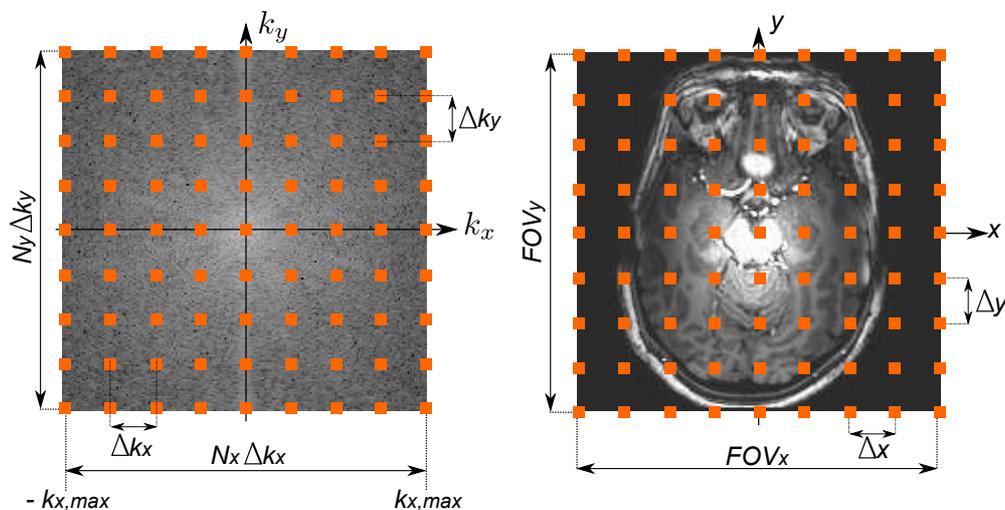


Figure 2.9: Relationship between k -space (left) and image-space (right). The image is first sampled in the k -space by applying magnetic field gradients. Frequency transform of the k -space samples is done to obtain the MR image. Adapted from [274].

structured image,

$$\Delta x = \frac{1}{2k_{x,max}}. \quad (2.22)$$

This most common way of sampling the k -space is intrinsically two-dimensional, with an image slice selectively sampled with successive RF excitation pulses. Nevertheless, a three-dimensional sampling can also be done. This is usually known as volume imaging [236, 328]. In volume imaging, the slice selective pulse is eliminated and 3D k -space trajectories are used to encode the spatial information in the third dimension. For instance, the most commonly used trajectory in volume imaging is a cartesian grid with gradient echo sampling in the x -axis and then phase encoding gradients along the y - and z -axis simultaneously. The advantage of 3D-acquisitions is a large improvement in the SNR of the image but the total acquisition time increases making the acquisition more vulnerable to motion artefacts [59].

2.2.4 Pulse Sequences

A pulse sequence is the temporal sequence of gradient pulses that is employed to navigate through the k -space so that an FID signal is collected at each sample of the k -space. In that sense, pulse sequence diagrams are particularly helpful to visualize how RF pulses and the magnitudes and durations of magnetic field gradients are manipulated over time to sample the entire k -space. Generally, a pulse sequence diagram displays a different line for each event of the sequence: the RF excitation

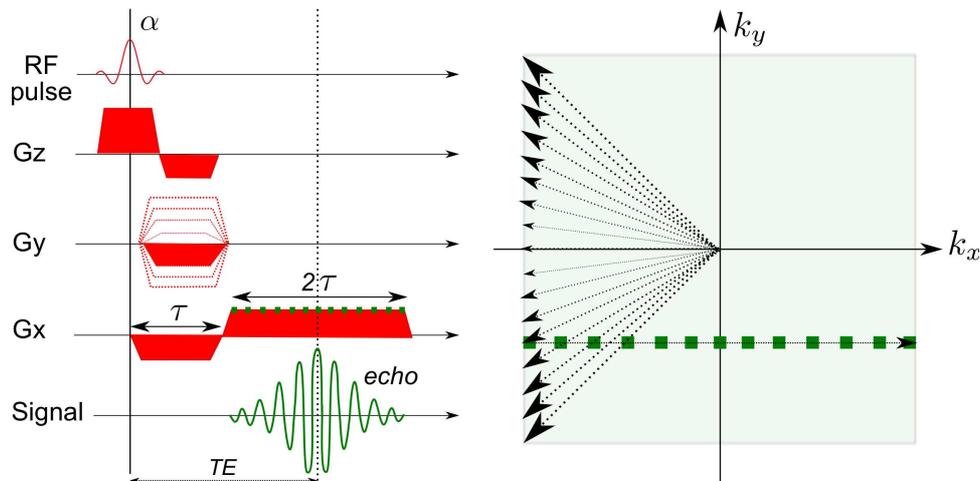


Figure 2.10: Gradient echo pulse sequence for one phase encoding step (left) and the corresponding sampling of one line in k -space (right). This pulse sequence diagram is successively repeated with changes in the polarity and amplitude of the phase encoding gradients to sample multiple lines in k -space. Adapted from [274].

pulse, the gradients in each direction (G_z or slice selection, G_x or frequency encoding, G_y or phase encoding) and the received FID signal.

An important parameter in MRI acquisition is the repetition time (TR), which is usually defined as the amount of time between two successive RF excitation pulses. For the single-shot EPI sequence usually employed in fMRI experiments, the TR corresponds to the time of two successive acquisition of the same slice, i.e. a complete volume.

Gradient Echo Imaging

As noted in §2.1.5, echoes can be generated by reverting the spin dephasing process. This can be done by applying a second RF refocusing pulse, producing spin echoes (SE), or by means of gradient fields, producing gradient echoes (GE). Spins across the sample precess at different frequencies due to linear gradient fields which makes the net magnetization dephase. If a gradient field with opposite amplitude would reverse the dephasing process, a gradient echo is produced. Figure 2.10 depicts the pulse sequence diagram and the k -space trajectory for the spin-warp sequence [106] with gradient echo imaging.

First, a frequency selective RF pulse is applied simultaneously with a slice selection gradient G_z to select a particular slice of the object. This RF pulse flips the bulk magnetization down by a flip angle α into the transverse plane. This is

followed by a reverse slice selection gradient to reverse the initial dephasing in the slice so that the spins at the top of the slice have the same phase as the spins at the bottom. After a period of time, which allows to modify the T_2^* contrast of the image, magnetic field gradients are applied in the x - and y -direction so that the sampling starts at the appropriate point in k -space. A frequency encoding gradient is first applied in reverse polarity during time τ to shift the magnetization phase to $-k_{x,max}$ (prephasing) at the same time as a phase encoding gradient to shift it to $k_{y,n}$. The purpose of the prephasing gradient is to prepare the transverse magnetization so that an echo is produced at a later time. Immediately afterwards, a frequency encoding gradient of the same amplitude is applied for 2τ while samples in the corresponding k_y -space line are recorded. Note that the maximum magnitude of the gradient echo occurs time τ after the start of the frequency encoding gradient, coinciding with the centre of the k_y -space line (see §2.1.5).

FLASH

The GE sequence is repeated after the net magnetization returns to equilibrium with a variable phase encoding gradient G_y in order to sample the k -space fully. This sequence is known as Fast-Low Angle Shot (FLASH) [156], and it is primarily used for fast structural or anatomical T_2^* -weighted imaging (see §4.3). The . Figure 2.11 shows the pulse sequence diagram and k -space trajectory of the FLASH sequence with 2D cartesian encoding in k_x and k_y . The sequence consists of multiple repetitions of the GE pulse sequence where a train of RF excitation pulses of flip angle α applied every repetition time (TR). Since the flip angle α is usually 20-30°, lower than using a 90° pulse, the time to recover the equilibrium magnetization is shorter and thereby a short TR can be used.

The RF pulse and slice selection gradient are applied simultaneously to generate transverse magnetization in the selected slice. After the RF pulse, the k -space is at zero phase. Then, a phase encoding gradient is applied before acquisition so that the phase of the signal is at the corresponding k_y . Next, a frequency encoding gradient in the k_x -direction is applied to acquire the corresponding k -space line, with initial reverse magnitude so that the gradient echo occurs at the centre of the k -space. Finally, a phase-rewinder gradient with reverse magnitude can be applied in the phase-encoding direction to return to zero-phase resulting in a balanced gradient

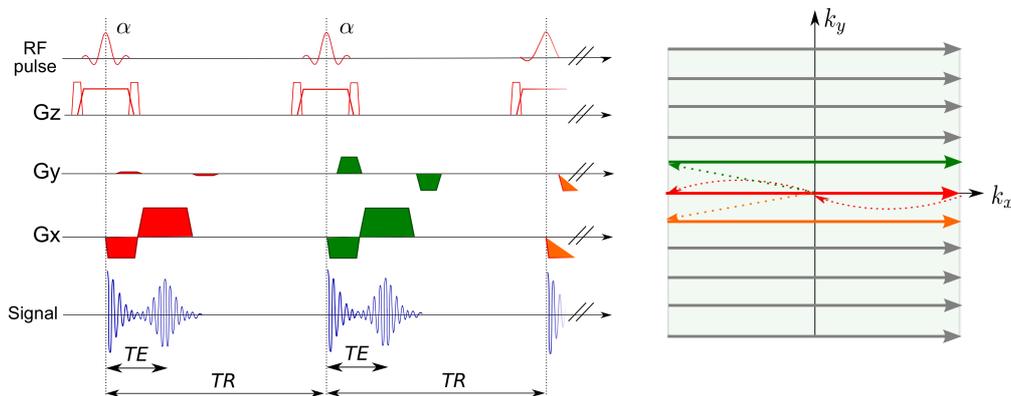


Figure 2.11: Two-dimensional version of the FLASH pulse sequence with GE (left) and the corresponding k -space trajectory (right). An RF pulse is applied every repetition time (TR) to sample one line in k -space. Adapted from [274].

waveform. In addition final spoilers are used to cancel any signal. If final spoilers are not used, the sequence is called balanced turbo field echo (bTFE) FLASH sequence. The process is repeated with variable magnitude in the phase-encoding gradients in order to generate more gradient echoes until the entire k -space is sampled. In multi-slice acquisition strategies, multiple slices are sampled for each TR period.

An important practical issue is how the longitudinal magnetization M_z behaves due to the first RF excitation pulses until achieving a steady-state magnetization [20, 178]. Assuming that the equilibrium magnetization due to the spin density population is M_0 and perfect spoiling of the transverse magnetization before each RF pulse, the longitudinal magnetization after each subsequent RF excitation pulse is given by

$$M_z(i+1) = M_0 \left(1 - e^{-TR/T_1}\right) + M_z(i) \cos(\alpha) e^{-TR/T_1}. \quad (2.23)$$

The steady-state is achieved when $M_z(i+1) = M_z(i)$ so that the steady-state longitudinal magnetization can be written as

$$M_z(i) = M_0 \frac{1 - e^{-TR/T_1}}{1 - \cos(\alpha) e^{-TR/T_1}}. \quad (2.24)$$

Therefore, a series of scans are usually left before starting to collect k -space samples in order to achieve a steady-state magnetization. This is particularly important in fMRI experiments since the first scans could greatly affect the results of posterior statistical analysis.

As noted above, each RF excitation pulse is followed by a GE acquisition. Therefore, similar to GE pulse sequences, the image is affected by T_2^* -relaxation effects, due to spin-spin interactions and field inhomogeneities (see §2.1.5). The FID signal after the RF pulse and the GE sequence can be written as [59, 20]

$$S_{GE} = M_0 \sin(\alpha) e^{-TE/T_2^*} \frac{1 - e^{-TR/T_1}}{1 - \cos(\alpha) e^{-TR/T_1}}, \quad (2.25)$$

where the maximum signal is obtained for a flip angle, α_E , called the Ernst angle [112], given by

$$\cos(\alpha_E) = e^{-TR/T_1}. \quad (2.26)$$

Therefore, one can control the amount of T_1 - and T_2^* -weighted imaging. For large flip angles, the signal is strongly dominated by T_1 -weighting. On the contrary, when the flip angle is smaller than α_E , the signal is relatively independent of T_1 and strongly T_2^* -weighted such that

$$S_{GE} = M_0 \sin(\alpha) e^{-TE/T_2^*}. \quad (2.27)$$

This expression is commonly used in GE fMRI experiments to demonstrate that the optimal signal change due to the BOLD effect, which depends on T_2^* -effect, is obtained when the signal is sampled at TE equal to T_2^* at rest (see §3.3).

In structural or anatomical T_2^* -weighted images, a relatively long TE and a TR of at least four or five times T_2 are used so that the transverse magnetization has enough time to decay to zero before the next excitation pulse. If $TR \ll 5T_2$ then an spoiler gradient must be used to destroy any remaining transverse magnetization. When low flip angles are used, the transverse magnetization is slightly perturbed whereas there is no noticeable change in the longitudinal magnetization.

Spin Echo Imaging

An alternative way of producing an echo is to apply a second RF inversion or refocusing pulse which rephases the spins, as it was illustrated in Figure 2.6. Figure 2.12 displays the pulse sequence diagram and the k -space trajectory for the spin echo (SE) sequence [157]. It can be observed that the main difference between the SE sequence and the GE sequence is the presence of a RF refocusing pulse instead

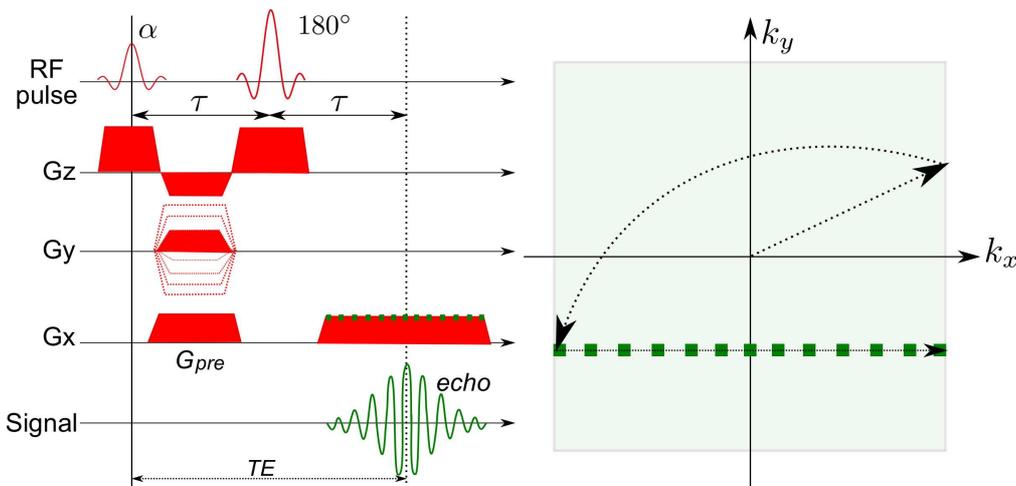


Figure 2.12: Spin echo pulse sequence for one phase encoding step (left) and the corresponding sampling of one line in k -space (right). The main difference with the GE sequence is the presence of a second RF pulse of 180° to produce the spin echo. Similar to the GE sequence, this pulse sequence diagram is successively repeated with changes in the polarity and amplitude of the phase encoding gradients to sample multiple lines in k -space. Adapted from [274].

of simply refocusing with gradients. The RF refocusing pulse rotates the frequency encoding axis by 180° as shown in the k -space trajectory. The use of spin echoes can revert the signal drop-out effects due to the static field B_0 -inhomogeneities so that the signal will only be affected by T_2 dephasing. Nevertheless, gradient echo imaging is unable to refocus field inhomogeneities.

Inversion recovery

A third widely used pulse sequence is called inversion recovery (IR). This sequence begins with an inversion 180° pulse, and then after a delay depending on the required contrast between samples or tissues, a regular SE or GE sequence begins with an excitation pulse. Contrary to the 90° RF excitation pulse, which nulls the longitudinal magnetization, the inversion pulse inverts the longitudinal magnetization in the $-z$ direction so that $M_z(0) = -M_0$. As shown in Figure 2.4, the longitudinal magnetization begins to recover with the spin-lattice relaxation process with an exponential decay with time constant T_1 . After a delay of TI an RF 90° excitation pulse is generally applied so that the longitudinal magnetization is flipped down into the transverse plane. Using (2.6), the magnitude of the longitudinal

magnetization at time TI is,

$$M_z(TI) = M_0 \left[1 - 2e^{-TI/T_1} \right]. \quad (2.28)$$

Additionally, this is affected by the T_2 or T_2^* dephasing effect in the transverse plane due to the spin-spin relaxation depending on whether a spin echo or gradient echo is used. If an additional 180° refocusing pulse at time $TE/2$ after the 90° excitation pulse, this produces a spin echo and the transverse decay depends on the larger T_2 , instead of T_2^* , due to the compensation of magnetic field inhomogeneities. The resulting signal for a SE IR sequence is given by [59, 20]

$$S_{IR} = M_0 e^{-TE/T_2} \left[1 - 2e^{-TI/T_1} + 2e^{-(TR-TE/2)/T_1} - e^{-TR/T_1} \right]. \quad (2.29)$$

If the TR of the image is much larger than TI , the signal is approximately

$$S_{IR} \propto M_0 e^{-TE/T_2} \left[1 - 2e^{-TI/T_1} \right]. \quad (2.30)$$

Note that if the inversion time TI is much longer than T_1 , the longitudinal magnetization is completely recovered, and the 180° inversion pulse does not produce any effect. But if TI is comparable to T_1 , the recovery is incomplete and the signal strongly depends on T_1 -relaxation producing a T_1 -weighted image. In particular, there is a null point when the longitudinal magnetization passes through $M_z = 0$ at time $TI = T_1 \ln 2$. If the 90° pulse is applied at this time, the signal from tissues with that particular T_1 is suppressed. This procedure is called null-point imaging [287] and it can be useful to help the segmentation of images into grey matter, white matter and cerebrospinal fluid and it is used in §3.6 .

Echo Planar Imaging

In 1977 Mansfield proposed a fast imaging technique which would revolutionize the way MRI was performed, the echo planar imaging (EPI) sequence [235]. In the simplest version of EPI the k -space is sampled after application of a single RF pulse, which is commonly referred to as single-shot EPI. After the RF excitation pulse, alternate gradient pulses in the frequency encoding and phase encoding directions yield a train of gradient echoes. In the simplest case, the k -space is sampled in a

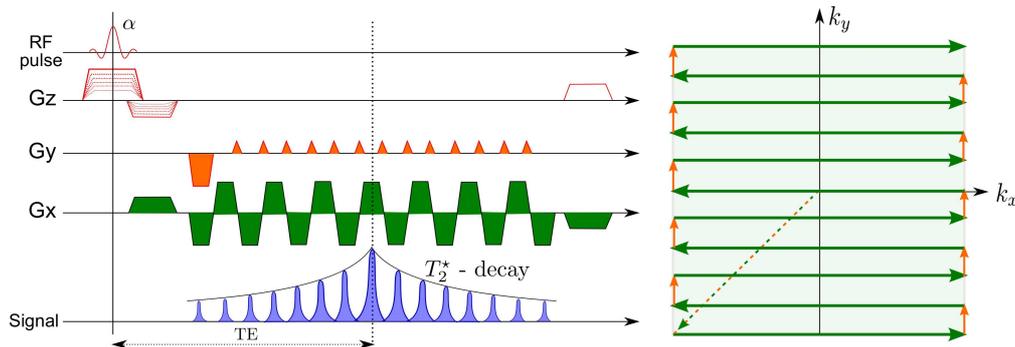


Figure 2.13: EPI pulse sequence (left) and the corresponding sampling of the k -space with the Modulus Blipped Echo-Planar Single-pulse technique (right). Adapted from [274].

zig-zag trajectory. Nowadays, the k -space is sampled in a rectilinear fashion with a Modulus Blipped Echo-Planar Single-pulse Technique [173, 271] (Figure 2.13), where a gradient in the frequency encoding direction is applied to shift the phase from side to side of the k -space followed by a smaller phase encoding gradient and so on. In addition, since the entire k -space is sampled after an unique RF excitation pulse, the number of sampled lines is lower than other pulse sequence techniques unless interleaved multishot EPI strategies are employed [57, 247].

As the echoes in EPI are formed from gradients, the EPI sequence is sensitive to B_0 -inhomogeneities unless an adequate shimming is employed [288] as well as susceptibility effects inducing geometric distortions artifacts or blurring in the image [183]. In addition, since the k -space trajectory is moving to the left in one line and back to right in the next line, small gradient imperfections can cause echoes be shifted in successive lines in k -space resulting in a ghost image superimposed over the actual image [59, 183]. Despite these drawbacks, the ability of acquiring 2D planar images in only a few milliseconds (fast-imaging) makes EPI a very popular sequence for functional MRI.

MPRAGE

The Magnetization Prepared Rapid Gradient Echo (MPRAGE) sequence [260] is a volume imaging technique which combines a periodic inversion pulse to enhance the T_1 -weighting in the image with a rapid GE acquisition. This allows the acquisition of high-resolution T_1 -weighted structural images with excellent contrast between gray and white matter (used in §4.3). Figure 2.14 shows the diagram for the MPRAGE pulse sequence. The sequence consists of three periods. For each slice acquisition

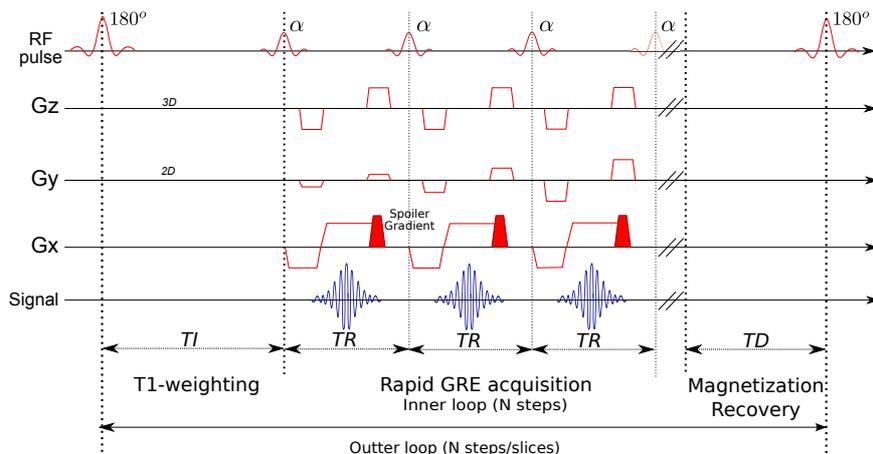


Figure 2.14: Pulse sequence diagram of the MPRAGE sequence. Adapted from [274].

(k_z -plane), an RF inversion pulse is first applied followed by a delay of time TI to generate a T_1 -weighting effect. This is followed by an acquisition stage with successive rapid GE sequences, usually using the FLASH sequence, to sample the image. Finally, the magnetization is allowed to recover for a time TD before the next inversion pulse is applied.

The excitation pulses produced by GE use low flip angles α so that the longitudinal magnetization does not considerably vary during the sampling of the x - y plane. In practice, since the repetition time between successive GE pulses is very short, a spoiled GE pulse is applied to destroy any remaining transverse magnetization before the sampling of the next k_y -line. In addition, since different lines of k_y are sampled with different longitudinal magnetization, there could exist variability of the signal intensity in the phase-encoding direction which results in a spatial smoothing of the image.

2.3 Conclusions

This chapter has introduced the basic principles of NMR and MRI. The NMR signal is created by the interaction of nuclear spins with applied magnetic fields at the corresponding resonance or Larmor frequency. Transverse and longitudinal relaxations control the amplitude of the NMR signal captured in MR scanners. Importantly, two different mechanisms to generate an NMR echo has been described: spin and gradient echoes. NMR principles are used in MRI to obtain images of a sample by applying spatially dependent magnetic fields and encode the properties of

the sample in the frequency or k -space. Several pulse sequences have been described to illustrate how k -space can be sampled and how MR images with different contrast can be produced. These techniques are relevant to understand the principles of functional MRI (fMRI) described in the following chapter.

Chapter 3

Functional Magnetic Resonance Imaging

This chapter describes the basis of functional magnetic resonance imaging (fMRI) based on the blood oxygenation level dependent (BOLD) endogenous contrast. The chapter begins with a short introduction to functional brain imaging techniques (§3.1). Subsequently, the chapter describes the physiological underpinnings of BOLD fMRI and the temporal characterization of the BOLD response (§3.2). The chapter also describes how to optimize sensitivity to detect BOLD signal changes in gradient echo imaging (§3.3), and the sources of noise in fMRI (§3.4). Next, the principles of two widely used fMRI data analysis methods, the general linear model and independent component analysis (§3.5), are also reviewed in this chapter. Based on these concepts, the chapter finishes with the description of two procedures to map variations of BOLD sensitivity across brain regions (BOLD sensitivity maps) and to correct test statistics depending on this variation (§3.6).

3.1 Functional Imaging

Neuroscientists and neurophysiologists have attempted to understand the functional organization of the brain for more than 100 years [305]. Functional imaging is the study of brain function based on the analysis of data acquired by means of neuroimaging techniques which are complementary in their temporal and spatial resolution (see Figure 3.1). The aim of these techniques is to resolve the neuro-

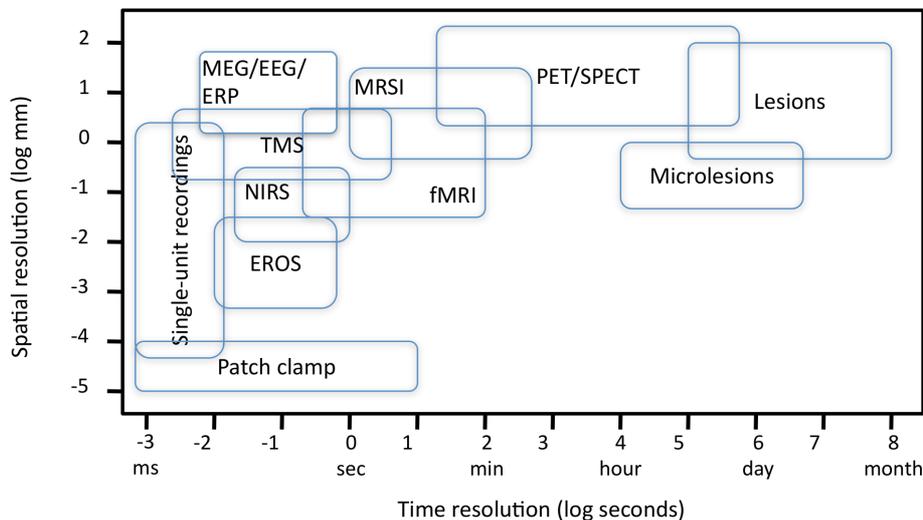


Figure 3.1: Spatial and temporal resolution of functional brain mapping techniques.

anatomical localization and dynamic changes of cortical activations.

Electrophysiological techniques such as electroencephalography (EEG) or magnetoencephalography (MEG) allow non-invasive recording of electric (EEG) or magnetic potentials (MEG) linked to the activity generated by an ensemble of neurons. They provide excellent temporal resolution (between 10 and 100 ms), but relatively poor spatial resolution (several mm or cm) [19, 159, 267]. This contrasts with the high spatial resolution achieved with invasive electrophysiological techniques, such as patch clamps [262], single-unit or multi-unit recordings [14] and electrocorticography (EcoG) [253, 266].

fMRI and Positron Emission Tomography (PET) techniques provide information about cerebral blood flow and oxygen metabolism indirectly associated with neural activations. Both techniques provide relatively high spatial resolution of the order of millimeters and are able to capture neuronal activations in deep brain structures, but exhibit limited temporal resolution due to the sluggish dynamics of hemodynamic changes. The main advantage of fMRI over PET is its higher temporal resolution and therefore is more suitable to study the temporal responses to short events. However, PET techniques have the advantage of measuring well-defined physiological quantities such as cerebral blood flow (CBF), cerebral blood volume (CBV) or cerebral metabolic rate of oxygen (CMRO₂) [126, 127].

MR spectroscopy imaging (MRSI) also enables metabolic imaging, providing

quantitative measurements of the relative concentrations of ions and compounds that contain a nucleus with a magnetic moment (e.g. ^1H , ^{31}P or ^{23}Na) within a brain region. The different nuclei composing a molecular structure can be resolved due to differences in their resonance frequency (measured in parts per million or ppm) in the proton magnetic resonance spectrum. However, the spatial and temporal resolution achieved with MRSI is lower than that achieved with fMRI techniques, even for proton MRSI [259].

Optical diffusion imaging techniques, such as near infrared spectroscopy (NIRS) [195, 355] or event-related optical signalling (EROS) [152, 153], also measure changes in the cortical blood flow by positioning an optical imaging device on the scalp and measuring the absorption or scattering of light. However, since the light must go through the skull, these techniques have lower spatial resolution than MR imaging techniques in addition to being restricted to the study of the cortical surface. The temporal resolution of EROS is of the order of milliseconds as with MEG and EEG, but that of NIRS is similar to fMRI since it studies blood flow. An advantage of optical imaging is that its instrumentation is relatively low cost and portable, compared to MRI or MEG.

Finally, transcranial magnetic stimulation (TMS) is a noninvasive technique that induces local currents by applying short but powerful magnetic fields close to the scalp. These currents can be used to stimulate neuronal activity in a particular region of the brain or to disrupt neuronal processing by creating a virtual brain lesion [276, 277].

3.2 BOLD Functional MRI

MRI techniques can be used to detect the increased metabolic demand associated with brain function based on the BOLD contrast, changes in CBV using contrast agents and in CBF using arterial spin labeling techniques (ASL) [59]. BOLD fMRI requires no exogenous contrast agent and is more sensitive than CBF contrast and so it is widely used for functional imaging since it was introduced by Owaga et al. in 1990 [268, 269].

In the following, some basic concepts related to the physical and physiological basis of the BOLD effect are reviewed in order to fully understand the assumptions

and methods presented in this thesis.

3.2.1 Physiological basis of the BOLD contrast

The signal contrast of BOLD fMRI images originates from changes in local magnetic susceptibility, χ , due to differences in the oxygen concentration of the haemoglobin in blood (see Figure 3.2). As local neuronal activity increases, there is an increase in local oxygen consumption which results in increased supply of oxygenated blood and this is delivered to the tissue by passive diffusion across the capillary blood vessels. When haemoglobin is bound to oxygen (oxyhaemoglobin), it is slightly diamagnetic compared to tissue so that the applied magnetic field is repelled by the molecule, while deoxygenated haemoglobin is paramagnetic compared to tissue and the applied magnetic field is pulled to the molecule [279]. Consequently, the magnetic field is distorted in the vicinity of deoxygenated red blood cells, creating higher local magnetic field gradients in the surrounding tissue leading to spin dephasing [338]. With less oxygenation, the spin dephasing effect is increased, shortening T_2^* of the tissue and reducing the amplitude of the MR signal in T_2^* -weighted images. Conversely, with higher oxygen concentration the susceptibility of the blood becomes similar to that of the surrounding tissue so the local magnetic field gradient is reduced, T_2^* increases and the measured MR signal amplitude also increases by a few percent. This is the principle of BOLD fMRI: changes in blood oxygenation are used as endogenous contrast in T_2^* -weighted images to reveal cortical areas with functional activity where the demand and supply of oxygen increase [22, 31, 203, 268, 269, 346]

3.2.2 Temporal characterization of the BOLD response: Haemodynamic Response Function

The BOLD effect does not directly reflect neuronal activity, but rather measures the haemodynamic response associated with it [222, 223]. The relationship between the hemodynamic response and the underlying neuronal activity is complex involving dynamic changes in CBF, CBV and CMRO₂ [59]. After neuronal activity increases in a brain region, there is an initial decrease in blood oxygenation due to oxygen consumption which might cause a small initial dip in the haemodynamic response following the first second after the activation [113, 250]. Although this initial dip

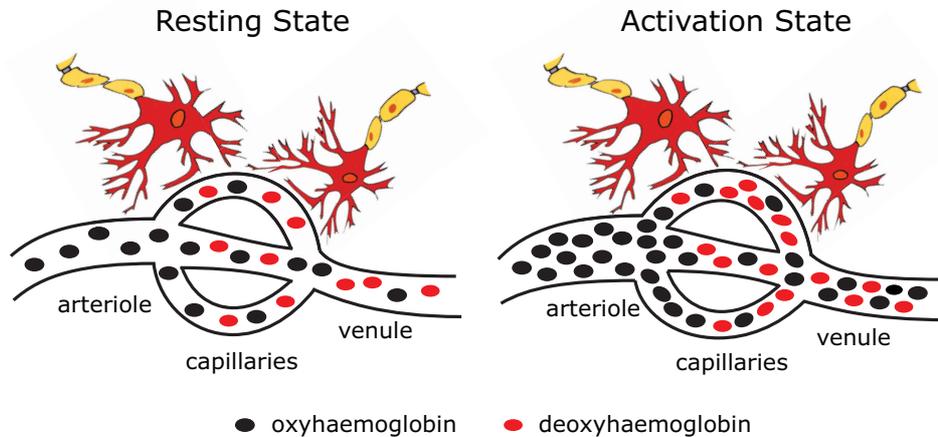


Figure 3.2: Blood Oxygenation Level Dependent Contrast. Changes in oxygen concentration in cerebral blood between activation and resting states causes a decrease in the local magnetic field gradient, increasing the MR signal measured with T_2^* -weighted images.

is not always observed in fMRI [28, 58, 174, 378], it is suggested that it maps more accurately the site of neural activity [104, 234]. Afterwards, the local supply of oxyhaemoglobin increases more than it is strictly demanded, probably to ensure a large oxygen gradient across capillary walls so that there is a high rate of transfer of oxygen or glucose to tissue [222], generating a positive BOLD response due to an excess of oxyhaemoglobin. Negative BOLD responses have also been observed associated with neuronal deactivations [319]. Regardless of the polarity of the response, the BOLD response peaks between 5 and 8 s after the activation starts and its amplitude depends on the type of stimulus and the magnetic field strength. For instance, for visual stimulation the signal change is 2-3 % at 1.5T, 4-6 % at 3T, 7-10 % at 7T [350]. After the stimulus ceases, there is a return of the BOLD response to baseline, often followed by a post-stimulus undershoot due to an increase of deoxyhemoglobin which may last for several seconds until the response returns to baseline. The cause of the post-stimulus undershoot is also not completely understood, whether this is a vascular, neural or metabolic effect [59, 75].

In summary, the time scale of the BOLD response is much slower than the time scale of neural activity and the return of the BOLD signal to baseline level after a short stimulus may occur more than 30 s from the stimulus onset. The temporal characteristics of the BOLD response are usually modelled by a haemodynamic response function (HRF). Figure 3.3 shows the shape of typical HRF, along with the initial dip for illustration of this effect. In this figure, the HRF plotted is the well-

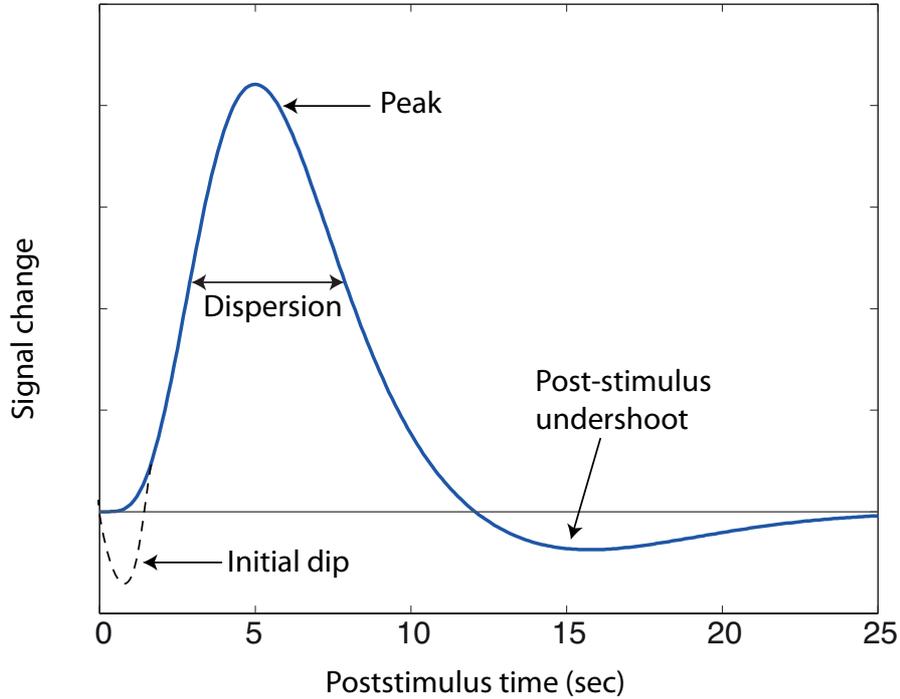


Figure 3.3: Temporal characterization of the haemodynamic response function. The shape in bold line corresponds to the canonical haemodynamic response function.

known canonical HRF [129, 130], which is defined as the difference of two-gamma functions :

$$h(t) = g(t; a_1, b_1) - \frac{1}{c}g(t; a_2, b_2), \quad (3.1)$$

where the Gamma function is given by

$$g(t; a, b) = \frac{b^a t^{a-1} e^{-bt}}{\Gamma(a)}. \quad (3.2)$$

The canonical HRF is parameterized by a time-to-peak (a/b) of 6 s and dispersion (a/b^2) of 1 s for the initial overshoot, and a time-to-peak of 16 s and dispersion of 1 s for the undershoot, with an overshoot-undershoot ratio (c) of approximately 6 [129]. Due to the variability of the BOLD response across trials, brain regions and subjects [4, 103, 162, 242, 323], other HRF shapes have been proposed in the literature such as single gamma functions [48, 79], two-gamma functions with different HRF parameters [143], Poisson functions [133] or Gaussian functions [199, 295].

3.3 BOLD sensitivity in Gradient Echo Echo Planar Imaging

The vast majority of fMRI experiments make use of the Echo Planar Imaging (EPI) pulse sequence [235] in order to acquire T_2 - or T_2^* -weighted 2D-images in less than 100 ms. Furthermore, Gradient Echo (GE) EPI is commonly used since it provides greater sensitivity to the BOLD contrast than Spin Echo (SE) EPI, although SE provides better spatial specificity of the location of the activation [376, 377]. In GE EPI, the T_2^* contrast in a voxel depends on microscopic field heterogeneity due to the presence of deoxyhaemoglobin, macroscopic field heterogeneity due to large vessels and magnetic field inhomogeneities, and the nuclear T_2 -relaxation rate (R_2). In total, the decay rate $R_2^* = 1/T_2^*$ can be expressed as

$$R_2^* = R_2 + R_2'. \quad (3.3)$$

With GE EPI the signal is acquired with longer TR in order to minimize T_1 -relaxation effects and at the optimal Ernst angle to maximize the amplitude of the signal [172]. In that regime, the MR signal is given by

$$S(t) = M_0 e^{-tR_2^*}, \quad (3.4)$$

and the BOLD signal due to changes in R_2^* is given by

$$\frac{\Delta S}{\Delta R_2^*} = -M_0 t e^{-tR_2^*}. \quad (3.5)$$

Figure 3.4 plots the magnitude of the signal change as a function of time for different simulated values of T_2^* at rest. The maximum BOLD contrast is achieved when the signal is sampled with an echo time TE equal to the value of T_2^* of the tissue in the region of interest. Hence, the BOLD contrast is given by

$$\begin{aligned} \Delta S &= -M_0 \Delta R_2^* TE e^{-TER_2^*} \\ &= -S \Delta R_2^* TE, \end{aligned} \quad (3.6)$$

where $S = M_0 e^{-TER_2^*}$ is the signal at rest, and $\Delta R_2^* = R_{2,a}^* - R_2^*$ is the change in

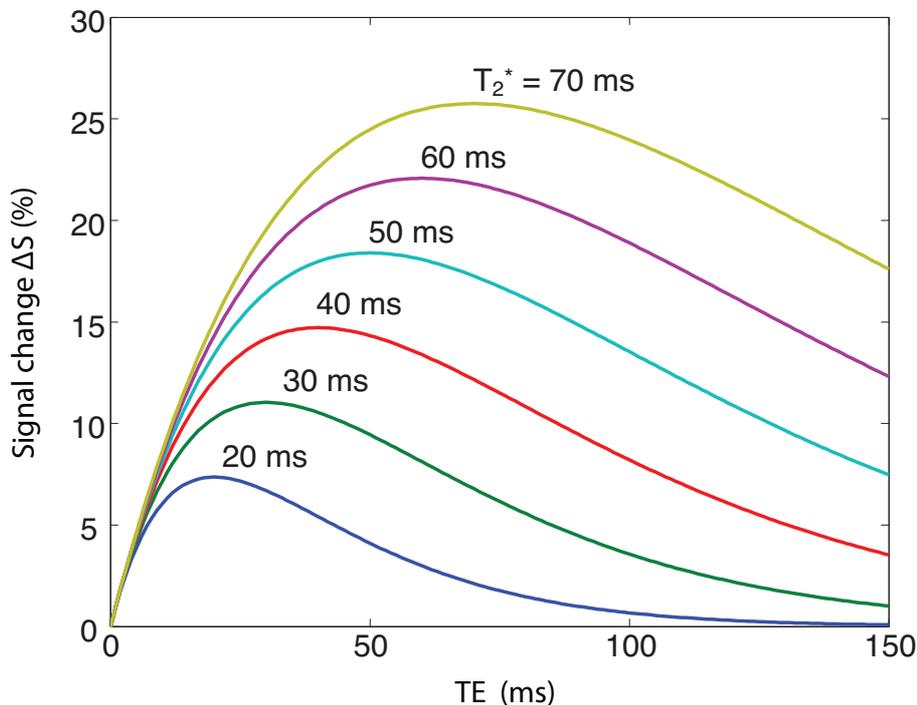


Figure 3.4: Magnitude of the BOLD contrast as a function of TE for different values of T_2^* . The maximum of the BOLD contrast occurs at $TE = T_2^*$.

R_2^* , where $R_{2,a}^*$ is the value of R_2^* due to the BOLD effect. Setting TE to a value shorter than T_2^* might also be appropriate in order to reduce the effect of susceptibility differences between voxels, whereas values longer than T_2^* are appropriate to minimize the venous component of the signal [283]. In practice, R_2^* or T_2^* maps can be obtained by fitting the linear decay of the T_2^* -weighted signal acquired with a multi-echo EPI sequence [184]. For instance, with a double-echo GE EPI sequence the signal is acquired at two echo times, one smaller and another larger than the average T_2^* value. If S_a and S_b are the signals acquired at the first echo time TE_a and second echo time TE_b respectively, the voxelwise R_2^* value can be estimated as:

$$R_2^* = \frac{\text{Ln}(S_a) - \text{Ln}(S_b)}{TE_b - TE_a}. \quad (3.7)$$

The BOLD sensitivity depends on the values of S , TE and R_2^* . Mapping the variability in BOLD sensitivity is necessary at any MR field strength to interpret the final fMRI results, but it becomes crucial at ultra-high magnetic fields, such as 7T, owing to the strong variations in R_2^* between brain regions due to B_0 -inhomogeneity effects and signal drop-offs in the signal magnitude S in peripheral regions of the

brain due to spatial B_1 -inhomogeneity effects [21, 95, 96, 361].

3.4 Noise in BOLD fMRI

The temporal signal-to-noise ratio, $tSNR$, of a voxel time series is defined as $tSNR = S/\sigma$, where S is average signal of the voxel time series, which depends on the baseline R_2^* , and σ is its standard deviation. In fMRI it is generally assumed that the standard deviation of the signal arises from noise-related fluctuations arising from at least four different sources [200, 201]:

$$\sigma^2 = \underbrace{\sigma_T^2 + \sigma_S^2}_{\sigma_0^2} + \underbrace{\sigma_B^2 + \sigma_{NB}^2}_{\sigma_P^2}. \quad (3.8)$$

In general terms, the term σ_0 is independent of the MR signal S , where σ_T represents the thermal noise from the subject and scanner electronics and σ_S describes other system noise and drift imperfections [105, 200]. On the other hand, σ_P includes the noise terms arising from physiological fluctuations, where σ_B includes the physiological noise arising from intrinsic brain fluctuations in $CMRO_2$, CBV and CBF that are proportional to the BOLD signal change ΔS and it can be expressed as $\sigma_B = c_1 \Delta S$, where c_1 is a constant, and σ_{NB} includes cardiac and respiratory fluctuations, and subject motion which also depend on the MR signal such that $\sigma_{NB} = c_2 S$. In short, the physiological noise term is signal-dependent such that $\sigma_P = \lambda S$, where

$$\lambda^2 = c_1^2 (\Delta R_2^*)^2 TE^2 + c_2^2. \quad (3.9)$$

Substituting (3.8) and (3.9), the $tSNR$ can be rewritten as:

$$tSNR = \frac{SNR_0}{\sqrt{1 + \lambda^2 SNR_0^2}}, \quad (3.10)$$

where $SNR_0 = S/\sigma_0$ is the signal-to-thermal noise ratio. As shown in Figure 3.5, the achievable $tSNR$ reaches an asymptotic limit of $1/\lambda$ as SNR_0 is increased. This asymptotic behaviour in $tSNR$ contrasts to an ideal scenario where no physiological noise exists ($\lambda = 0$). As this asymptotic limit is approached, further increases in SNR_0 due to higher MR fields or improvement in the RF system only translates into small improvements in $tSNR$.

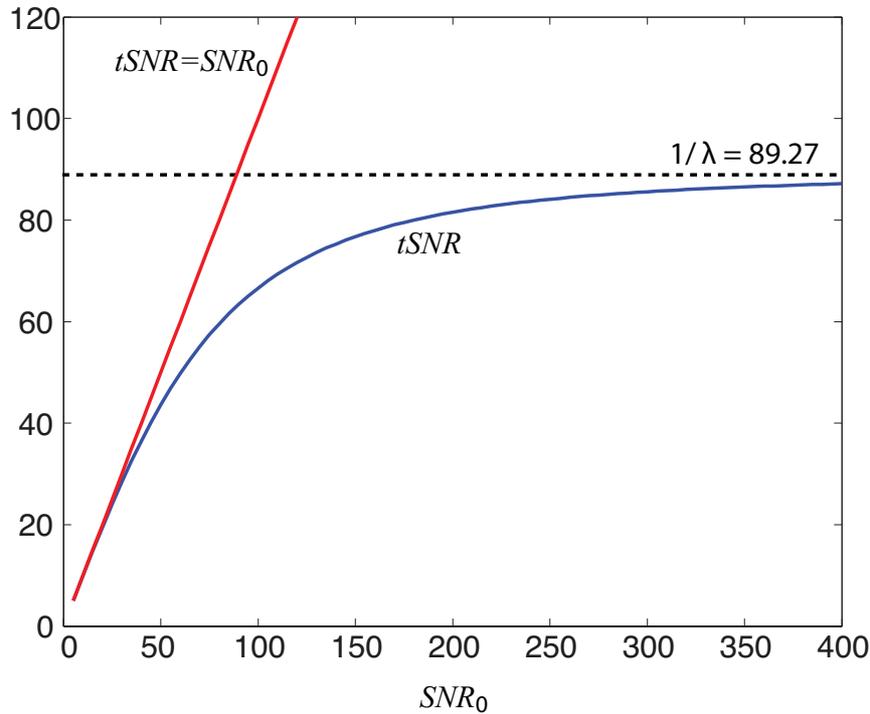


Figure 3.5: Temporal signal-to-noise ratio ($tSNR$) as a function of the signal-to-thermal noise ratio (SNR_0). There is an asymptotic limit in $tSNR$ equal to $1/\lambda$ due to the physiological noise terms depending on the MR signal. This contrasts with the linear increase in $tSNR = SNR_0$ when no physiological noise exist, i.e. $\lambda = 0$. The $tSNR$ curve is computed for $\lambda=0.0112$, which is the mean value of λ in GM tissue at 1.5, 3T and 7T [343].

The value of λ is voxel specific, with higher values in cortical gray matter (GM) than in white matter (WM) [200, 343]. The ratio of physiological to thermal noise (σ_P/σ_0) depends on the MR acquisition parameters such as flip angle, voxel resolution, coil type or parallel imaging strategy [44, 342, 343]. More specifically, thermal noise is dominant over physiological noise at 1.5T until approximately voxel volume of 70 mm^3 , whereas it is the opposite at 7T where the physiological noise is higher than thermal noise except in a very high resolution regime [44, 343]. Therefore, it is in fMRI data acquired at high spatial resolution, around 3 mm^3 , where the advantage of 7T MR scanners becomes really apparent since further increases in SNR_0 translate into increases in $tSNR$ [329]. This differentiation between thermal and physiological noise will be used in Chapters 5 and 6 to create realistic simulations of fMRI voxel time series.

3.5 fMRI data analysis

fMRI data analysis techniques aim to detect and characterize the BOLD response in order to shed light on the dynamics of the underlying neuronal activity. There exist many statistical methods to study BOLD fMRI data and therefore the choice of the statistical method is a critical step which determines the type of information that the results of the analysis provide. Rather than providing a complete review (e.g., see [207, 311]), we will focus on model-based analysis based on a general linear model (GLM) and model-free analysis based on independent component analysis (ICA). Both GLM and ICA will serve as benchmarks with which compare the methods introduced in this thesis.

3.5.1 General Linear Model

The general linear model (GLM) is the most widespread method for analyzing fMRI data and it encompasses many of the principles used by more complex techniques [129, 132, 133, 372]. In GLM analysis, the fMRI data observed in a voxel k at time t is mathematically described as [133]

$$y_t^k = x_{t1}^k \beta_1 + \dots + x_{tl}^k \beta_l + \dots + x_{tL}^k \beta_L + e_t^k, \quad (3.11)$$

where x_{tl} and β_l , for $l = 1, \dots, L$, are the explanatory variables at time t and their corresponding unknown weights, respectively. Having collected N observations for $t = 1, \dots, N$, the fMRI time course of the voxel can be written as:

$$\begin{aligned} y_1^k &= x_{11}^k \beta_1 + \dots + x_{1l}^k \beta_l + \dots + x_{1L}^k \beta_L + e_1^k, \\ &\vdots \\ y_t^k &= x_{t1}^k \beta_1 + \dots + x_{tl}^k \beta_l + \dots + x_{tL}^k \beta_L + e_t^k, \\ &\vdots \\ y_N^k &= x_{N1}^k \beta_1 + \dots + x_{Nl}^k \beta_l + \dots + x_{NL}^k \beta_L + e_N^k \end{aligned} \quad (3.12)$$

Let $\mathbf{y}_k = (y_1, \dots, y_N)^T$, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_L)^T$ and $\mathbf{e}_k = (e_1, \dots, e_N)^T$ be column vectors denoting observations, weights and noise terms of voxel k , and $\mathbf{X}_k = [\mathbf{x}_1, \dots, \mathbf{x}_L]$ as the $N \times L$ design matrix with regressors $\mathbf{x}_l = (x_{1l}, \dots, x_{tl}, \dots, x_{Nl})^T$. The entire

linear model can be written as

$$\mathbf{y}_k = \mathbf{X}_k \boldsymbol{\beta}_k + \mathbf{e}_k. \quad (3.13)$$

The design matrix comprises those fixed effects that may explain some variability of the fMRI time series, such as experimental regressors that model signal changes for each experimental condition [132, 133]. Apart from experimental regressors, it is usual to include regressors accounting for artifacts apparent in the fMRI signal due to subject motion, such as the realignment parameters computed during registration of the MRI scans to a common reference image [137], or regressors accounting for systematic drifts or low frequency trends [229]. The standard way of analysis in fMRI is to define the same design matrix for all the voxels. However, this is not a necessary condition and different design matrices could be proposed for different voxels in order to better capture the variability of the fMRI time course across different voxels [230, 297]. This will be the case for the Paradigm Free Mapping techniques described in the following chapters. On the other hand, the noise term includes all random effects which vary every time the data is scanned again. Based on the central limit theorem, the noise term in fMRI data can be well approximated by a Multivariate Normal distribution with zero mean and covariance matrix $\sigma^2 \boldsymbol{\Sigma}$, i.e. $\mathbf{e}_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{\Sigma})$.

The experimental regressors included in the design matrix attempt to model the BOLD response to a stimulus. Assuming a linear time-invariant system, the BOLD response can be modelled as the linear convolution of the HRF, $h(t)$, and a stimulus signal, $s(t)$:

$$x(t) = h(t) * s(t) = \int_0^\infty h(\tau) s(t - \tau) d\tau, \quad (3.14)$$

where $*$ denotes the convolution operator [270]. Figure 3.6 shows the BOLD response to different stimuli of increasing duration assuming a canonical HRF. Depending on the goals of the experiment, the duration of the stimuli can be short, usually less than 5 or 10 s, generating an event-related response, which is optimal to estimate the temporal characteristics of the HRF. On the other hand, if the experimental paradigm comprises an alternation of long, sustained stimuli, it is known as block design or epoch-related, which is optimal in terms of maximizing the detection of the BOLD response [219].

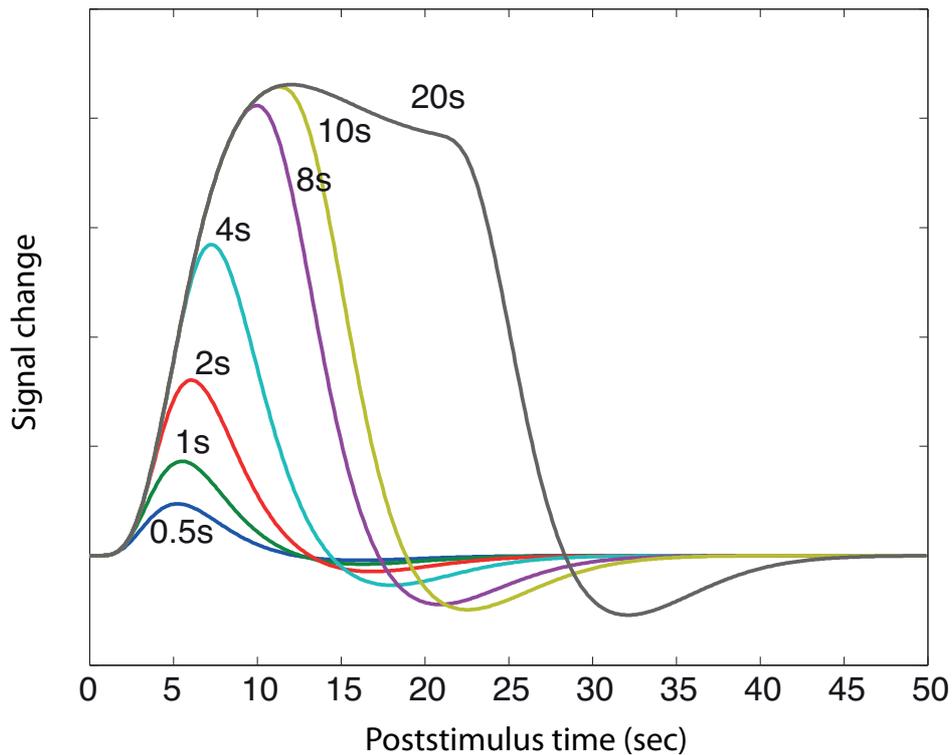


Figure 3.6: Model of the BOLD response generated by linear convolution of stimulus with increasing duration with the canonical HRF (Adapted from [129]).

In many experiments, the researcher is interested in the contrast between two different tasks or stimuli, e.g. task A and task B (Figure 3.7). In practice, several experimental regressors are usually created for each task by the convolution of the stimulus signal, $s(t)$, with a set of basis functions characterizing the shape of the HRF, rather than using a single HRF in order to account for the HRF variability [129, 130]. This helps to enhance the sensitivity of the analysis to detect an effect providing the basis set is properly designed. Otherwise, increasing the number of regressors in the design matrix by using more basis functions increases the variability of the model estimates and reduces the statistical significance of the detected response [214]. A common set of basis functions is the informed basis set, which includes the canonical HRF and the temporal and dispersion derivatives (Figure 3.8) [130]. The informed basis set makes possible to test for temporal shifts in the delay and dispersion of the canonical HRF [374]. Other common sets of basis functions used in fMRI analysis are the Fourier or the Finite Impulse Response (FIR) basis sets, which do not assume any a-priori shape of the HRF [129]. Finally, after convolution, each regressor is sampled at the temporal resolution of the fMRI time series,

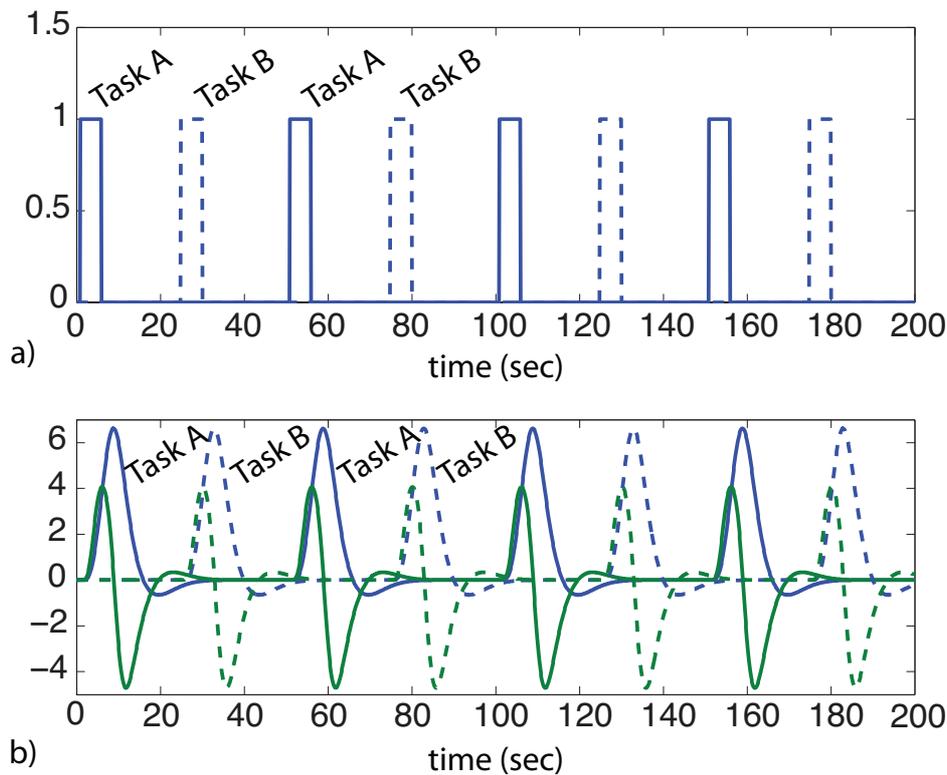


Figure 3.7: Creation of regressors for the design matrix. a) The signals corresponding to two different stimuli, task A and task B, b) are convolved with the set of basis functions, in this case the canonical HRF (blue) and its first temporal derivative (green), in order to create the predicted BOLD regressors (Adapted from [184]).

equal to the repetition time (TR).

It is very important to note that the assumption of linearity in the BOLD response is an oversimplification which is useful to allow us formulate tractable mathematical models of the BOLD response. In particular, the coupling between stimuli and BOLD responses might exhibit nonlinear effects under certain stimulus patterns, such that the amplitude of the response to a sustained stimulus is not as high as it could be predicted from a short stimulus [39, 48, 143, 285, 327, 353] or in rapid event-related designs with short overlapping responses in contrast to spaced designs [92, 360]. A practical solution to this is to model the nonlinear effects in the GLM formulation with regressors which consider the quadratic (second order Volterra kernels) or cubic (third-order Volterra kernels) interactions between the linear regressors [134].

Once the design matrix for a given experiment is defined, estimates of the weights β can be computed by means of classical maximum likelihood (ML) or Bayesian

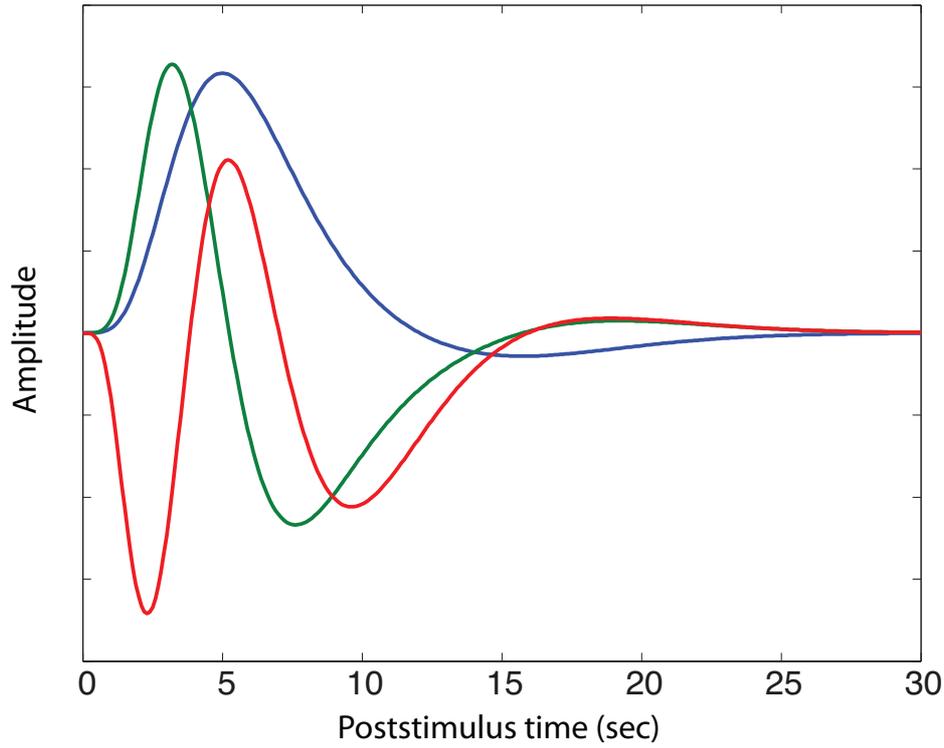


Figure 3.8: Informed basis set: Canonical HRF (blue) and its temporal (green) and dispersion (red) derivatives (Adapted from [129])

estimators [181]. Adopting a classical inference approach and assuming that the covariance matrix of the errors is known, the best linear unbiased estimator of β is the generalized least squares estimator:

$$\hat{\beta} = \arg \max_{\beta} (\mathbf{y} - \mathbf{X}\beta)^T \Sigma^{-1} (\mathbf{y} - \mathbf{X}\beta) \quad (3.15)$$

where

$$\hat{\beta} = (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{X}^T \Sigma^{-1} \mathbf{y}. \quad (3.16)$$

In practice, the covariance matrix of the errors Σ is not known and must be estimated from the data. An iterative reweighted least squares procedure is then adopted which first fits a model assuming uncorrelated errors, i.e. $\Sigma = \mathbf{I}$, such that β is estimated with the ordinary least squares estimator:

$$\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y} = \mathbf{X}^+ \mathbf{y}, \quad (3.17)$$

where $\mathbf{X}^+ = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$ is the Moore-Penrose pseudoinverse of the matrix \mathbf{X} . Then, the residuals of the uncorrelated-error model are computed as $\mathbf{r} = \mathbf{y} - \hat{\mathbf{y}}$,

where $\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}}$. Substituting Eq. (3.17), the residuals can be computed as $\mathbf{r} = \mathbf{R}\mathbf{y}$, with the residual projection matrix defined as $\mathbf{R} = \mathbf{I} - \mathbf{X}\mathbf{X}^+$. Next, the covariance matrix $\hat{\boldsymbol{\Sigma}}$ required to compute the GLS estimates in Eq. (3.16) can be estimated from these residuals. This process can be repeated iteratively to estimate $\hat{\boldsymbol{\beta}}$ and $\hat{\boldsymbol{\Sigma}}$, but in practice only two iterations are sufficient to achieve efficient estimates [373]. Importantly, estimating the $N(N - 1)/2$ correlation coefficients of the matrix $\hat{\boldsymbol{\Sigma}}$ in a voxel-by-voxel basis results in biased estimates at large correlation lags. Consequently, the temporal correlation structure is usually modelled by an autoregressive (AR) or autoregressive moving average (ARMA) stochastic process.

It can be shown that the variance of the GLS estimates is given by

$$\text{var}(\hat{\boldsymbol{\beta}}) = \sigma^2(\mathbf{X}^T\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}. \quad (3.18)$$

Notice that the model of the temporal correlation of the error becomes decisive when estimating the standard errors of the estimates and assessing the significance of an effect, along with the error variance estimate. Assuming that $\boldsymbol{\Sigma}$ is known or estimated, the error variance estimate $\hat{\sigma}^2$ is computed from the residuals of the model $\mathbf{r} = \mathbf{R}\mathbf{y}$ as

$$\hat{\sigma}^2 = \frac{\mathbf{r}^T\mathbf{r}}{\text{tr}(\mathbf{R}\boldsymbol{\Sigma})}. \quad (3.19)$$

Based on the Satterthwaite approximation [312], the effective degrees of freedom of this estimate can be approximated by

$$\nu = \frac{2E(\hat{\sigma}^2)^2}{\text{var}(\hat{\sigma}^2)} = \frac{\text{trace}(\mathbf{R}\boldsymbol{\Sigma})^2}{\text{trace}(\mathbf{R}\boldsymbol{\Sigma}\mathbf{R}\boldsymbol{\Sigma})}, \quad (3.20)$$

and the distribution of $\nu\hat{\sigma}^2/\sigma^2$ is well described by a χ^2 distribution with ν degrees of freedom [129, 184].

In classical statistical inference, the significance of the effects of the GLM regressors related to a given stimulus is assessed by means of T -tests and F -tests. The question to evaluate with a T -test is about a linear combination of effects, $\mathbf{c}^T\boldsymbol{\beta}$, defined with a contrast vector \mathbf{c} . For instance, to test whether the first regressor in the design matrix \mathbf{X} is relevant against a null condition, the contrast vector is defined as $\mathbf{c} = (1 \ 0 \ \dots \ 0)^T$ so that $\mathbf{c}^T\hat{\boldsymbol{\beta}} = \hat{\beta}_1$, or to test whether the difference between the first and second regressors is significant, the contrast vector is $\mathbf{c} = (1 \ -1 \ \dots \ 0)^T$

so that $\mathbf{c}^T \hat{\boldsymbol{\beta}} = \hat{\beta}_1 - \hat{\beta}_2$. In general, given a column vector c with L weights, the distribution of $\mathbf{c}^T \hat{\boldsymbol{\beta}}$ is given by

$$p(\mathbf{c}^T \hat{\boldsymbol{\beta}} \mid \sigma) \sim \mathcal{N}(\mathbf{c}^T \boldsymbol{\beta}, \sigma^2 \mathbf{c}^T (\mathbf{X}^T \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1} \mathbf{c}), \quad (3.21)$$

conditional on σ . After accumulating the uncertainty in σ , $p(\mathbf{c}^T \hat{\boldsymbol{\beta}})$ is then seen to follow a Student's t -distribution with ν degrees of freedom, t_ν . A t -statistics can then be defined as

$$T(\mathbf{c}^T \hat{\boldsymbol{\beta}}) = \frac{\mathbf{c}^T \hat{\boldsymbol{\beta}}}{\hat{\sigma} \sqrt{\mathbf{c}^T (\mathbf{X}^T \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1} \mathbf{c}}}, \quad (3.22)$$

to test the null hypothesis $\mathcal{H}_0 : \mathbf{c}^T \hat{\boldsymbol{\beta}} = 0$ [129, 184].

F -tests are used to assess whether a group of regressors (or effects) is significant over the complete model. Assuming that the full model \mathbf{X} can be partitioned into two parts such that $\mathbf{X} = [\mathbf{X}_1 : \mathbf{X}_2]$, where \mathbf{X}_1 includes the effects to test and \mathbf{X}_2 is the reduced model, the parameters can also be bi-partitioned as $\boldsymbol{\beta} = [\boldsymbol{\beta}_1^T \boldsymbol{\beta}_2^T]^T$. After estimating $\hat{\boldsymbol{\beta}}$, we can compute the residual sum of squares of the full model $S(\boldsymbol{\beta}) = \mathbf{r}^T \mathbf{r}$, and the residual sum of squares of the reduced model $S(\boldsymbol{\beta}_2) = \mathbf{r}_2^T \mathbf{r}_2$, where $\mathbf{r}_2 = \mathbf{y} - \mathbf{X}_2 \hat{\boldsymbol{\beta}}_2$. Testing for the hypothesis that the regressors in $\hat{\boldsymbol{\beta}}_1$ are not significant, i.e. $\mathcal{H} : \hat{\boldsymbol{\beta}}_1 = 0$, is done with the following F -statistic:

$$F = \frac{S(\boldsymbol{\beta}_2) - S(\boldsymbol{\beta})}{S(\boldsymbol{\beta})} \cdot \frac{N - p}{p - p_2}, \quad (3.23)$$

where $p = \text{rank}(\mathbf{X})$ and $p_2 = \text{rank}(\mathbf{X}_2)$. This statistic follows an F -distribution, F_{d_1, d_2} , with $d_1 = p - p_2$ and $d_2 = N - p$ degrees of freedom. It can be shown that when \mathbf{X}_1 only includes one regressor, then the F -statistic is equal to the squared T -statistic, T^2 , for that regressor.

Significance p -values are computed from the T - and F -statistics with standard statistical procedures from the cumulative probability functions of the Student's t -distribution and F -distribution, respectively. The p -values are usually thresholded in order to decide, at a given level of significance, at which voxels the null hypothesis is rejected, creating a Statistical Parametric Map (SPM). Note, however, that we face a multiple hypothesis testing scenario where N voxels are tested. If rejection is done at a p -value of α , then αN voxels will be rejected even if the signal is random noise, where the null hypothesis is true, becoming false positives. Owing to the

multiple comparison problem, p -values must be corrected with a statistical correction procedure, such as Bonferroni, Gaussian Random Fields or False Discovery Rate [129, 141, 265, 370].

3.5.2 Independent Component Analysis

Independent Component Analysis (ICA) is used in fMRI as a model-free approach where no specific model for the fMRI data is assumed. Since the first papers applying ICA for fMRI data analysis [244, 245, 246], this technique has acquired an enormous popularity in the neuroscience community so that nowadays it is probably the most employed model-free technique [65]. Model-free or data-driven approaches are very useful in situations where the brain response to a given task may not be well-modelled, such as transient task-related signals or decaying responses during long-epoch stimuli.

Consider that the fMRI data is generated by the linear mixture of L components or sources, $\mathbf{X} = [\mathbf{x}_1^T \mathbf{x}_2^T \dots \mathbf{x}_L^T]^T$,

$$\mathbf{y}_k = \mathbf{X}\mathbf{a}_k, \quad (3.24)$$

where the vector $\mathbf{a}_k = [a_{1k} \dots a_{Lk}]^T$ represents the contribution of each source to the time series of voxel k , \mathbf{y}_k . For instance, the sources of the data can be task-related, motion-related, physiological signals and scanner drifts. Gathering the time courses of K voxels, we can write the linear mixing model

$$\mathbf{Y} = \mathbf{X}\mathbf{A}, \quad (3.25)$$

where $\mathbf{Y} = [\mathbf{y}_1^T \dots \mathbf{y}_K^T]^T$ is a $N \times K$ matrix including the voxel time courses, and $\mathbf{A} = [\mathbf{a}_1^T \dots \mathbf{a}_K^T]^T$ is a mixing matrix comprising the weights that specify the contribution of each source to each voxel, i.e. a spatial map of weights. The purpose of ICA is to blindly estimate the matrices \mathbf{A} and \mathbf{X} by assuming independence between the spatial maps of weights as in spatial ICA, or the source time courses as in temporal ICA, eventually decomposing the data into the original sources and so this problem is also known as blind source separation.

Figure 3.9 illustrates the differences between spatial and temporal ICA. On the

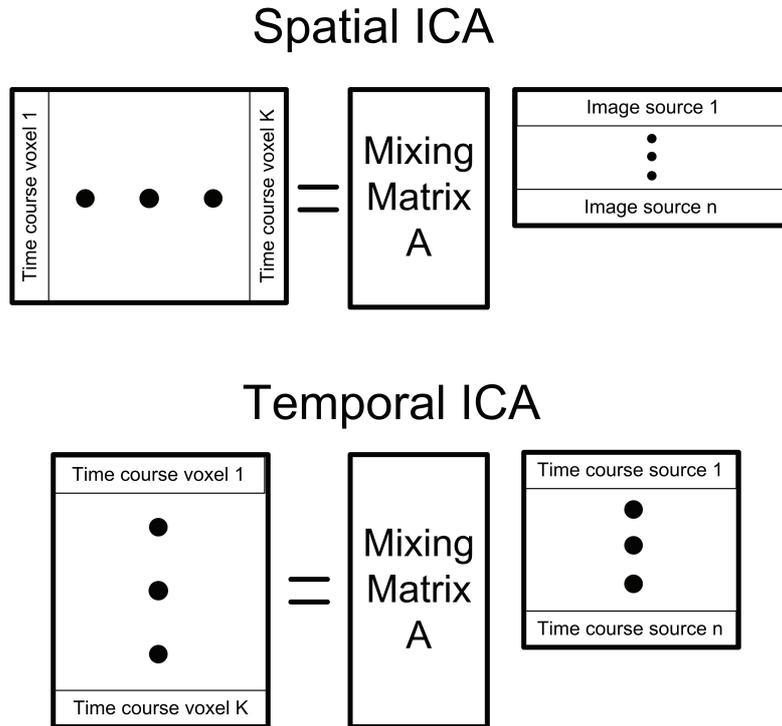


Figure 3.9: Spatial and temporal ICA.

one hand, the aim in spatial ICA (sICA) is to compute the unmixing matrix \mathbf{W} such that

$$\mathbf{A} = \mathbf{W}\mathbf{Y}. \quad (3.26)$$

The criteria to obtain \mathbf{W} is to make the spatial maps of weights, i.e. the rows of \mathbf{A} , statistically independent. Additionally, the algorithm also computes estimates of the source time courses as the columns of the pseudo-inverse of \mathbf{W} , \mathbf{W}^+ . On the other hand, one could consider a temporal ICA (tICA) model [66], where the signal matrix is \mathbf{Y}^T of size $K \times N$, i.e. $\mathbf{Y}^T = \mathbf{A}^T\mathbf{X}^T$, and the goal is to estimate the matrix $\widetilde{\mathbf{W}}$ such that

$$\mathbf{X}^T = \widetilde{\mathbf{W}}\mathbf{Y}^T, \quad (3.27)$$

and the sources time courses, i.e. the columns of \mathbf{X} , are statistically independent. In this case, the algorithm estimates the spatial maps of weights as the rows of $\widetilde{\mathbf{W}}^+$.

ICA-based approaches can be classified according to the criteria followed to compute the independent components, such as maximization of information transfer or maximum likelihood (Infomax) [29, 244, 245, 246], maximization of non-gaussianity (FastICA) [177, 27], mutual information minimization or entropy maximization

[111, 110], or joint approximate diagonalization of eigenmatrices (JADE) [69]. Importantly, ICA results in fMRI data depend on the algorithmic criteria [81, 114], the number of components to estimate [1, 212] or the initialization of the algorithm [168]. In practice, the dimensionality of the data is reduced by applying principal component analysis (PCA) in combination with information theoretic criteria prior to the ICA decomposition, which decreases the number of ICs to estimate. Among the algorithmic criteria, Infomax has demonstrated more consistent results than FastICA and JADE. However, since most ICA algorithms perform a stochastic search of the components, it is advisable to perform some type of bootstrapping procedure, such as ICASSO [168], to validate the components. Typically, the sICA formulation is adopted for fMRI analysis where the output is a set of spatial independent maps showing nonoverlapping brain networks with an associated temporal signature (time course). The argument in sICA is that each spatial component showing plausible haemodynamic time courses identifies a distinctive functional brain network with temporal coherence [13, 88]. Nevertheless, it has been argued that independence might not be the right mathematical framework for blind source separation in fMRI data, but that the decomposition must be driven by the sparsity of the components or their intersections, i.e. separability of the components [89].

Finally, in addition to detecting task-related signals, ICA has also been employed for data denoising as a preprocessing step. Artfactual components, such as motion-effects or physiological noise, are identified with ICA as they are assumed to be independent to the BOLD-related components, and then removed from the data before performing a GLM analysis [243, 282].

3.6 Correcting BOLD sensitivity maps

As noted in §3.3, the sensitivity to detect the BOLD response depends on physiological and physical parameters of the fMRI acquisition, i.e. S , TE and R_2^* , which vary across brain regions. In addition, the detection of the effect with a t -statistic also depends on the assumed model of the BOLD response. Combining both concepts, this section describes two approaches to map the required change in R_2^* , ΔR_2^* , so that the BOLD response is statistically significant with a certain p -value and viceversa; and compute correction factors for GLM-based statistics so

that a constant ΔR_2^* were measured uniformly regardless of the R_2^* variability, i.e. assuming that all voxels were acquired at their estimated T_2^* by virtually adapting TE .

3.6.1 Theory

A BOLD sensitivity map is defined in this study as a map showing on a voxel-by-voxel basis the ΔR_2^* required to detect statistical significant activation with a required p -value given the underlying variation in the baseline R_2^* .

From Eq. (3.22), the necessary signal change ΔS to reject the null hypothesis of non activation is proportional to the standard deviation of the noise σ and a statistical factor K ,

$$\Delta S = \sigma K. \quad (3.28)$$

Assuming uncorrelated noise, the factor K is defined from Eq. (3.22) as

$$K = T_{th} \sqrt{\mathbf{c}^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{c}}, \quad (3.29)$$

so that it depends on the design matrix \mathbf{X} (e.g. number of scans, number and duration of the events, the shape of the HRF) and the statistical significance required to accept the alternative hypothesis of activation T_{th} (or related p_{th}). Figure 3.10 plots curves of the factor K versus the scan duration in order to detect a single-trial event with statistical significance of $p = 10^{-6}$, for different stimulus duration (1, 2 and 4 s) and TRs (1, 2 and 4 s), where the linear model includes a unique regressor which is the convolution of the stimulus time course with the canonical HRF. For instance, the factor K to detect a single trial event of duration 2 s in a time series of duration 120 s with TR of 2 s, i.e. $N = 60$ scans, is 8.78 (shown with a black dot in Figure 3.10).

Combining (3.6) and (3.28), the ΔR_2^* required to significantly reject the non-activation hypothesis is given by

$$\Delta R_2^* = -\frac{K}{tSNR TE}. \quad (3.30)$$

Therefore, $tSNR$ is demonstrated to be the fundamental measure to map variations in BOLD sensitivity across brain regions [261, 275]. Since $tSNR$ depends on

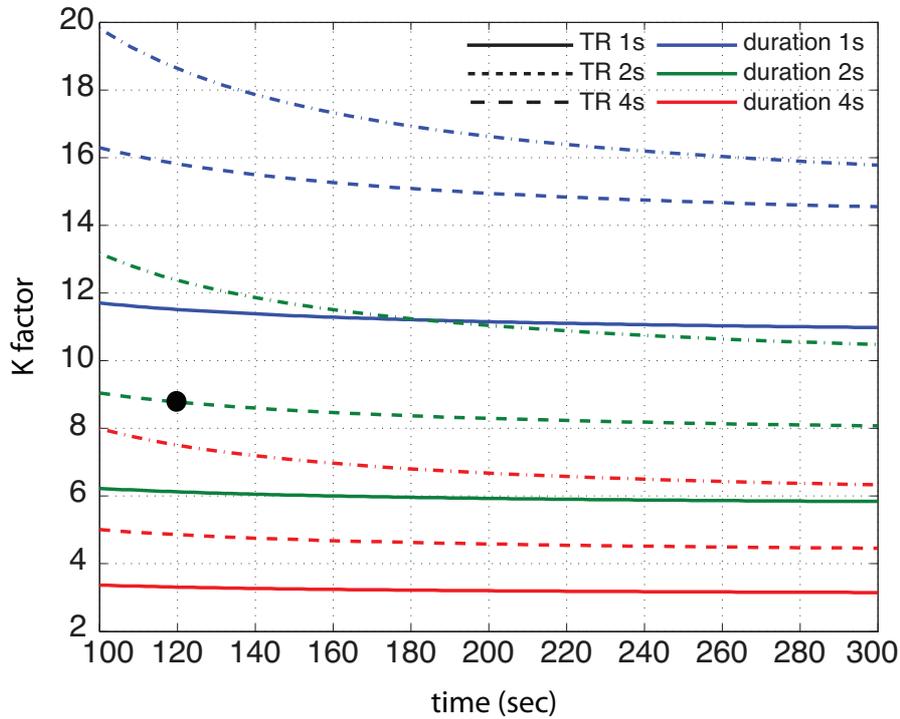


Figure 3.10: Statistical factor K in Eq. (3.28) in order to detect a single trial event of duration 1, 2 and 4 s with statistical significance of $p = 10^{-6}$ versus the scanning time when TR is 1, 2 and 4s. The dot point is the value of $K=8.78$ used in order to compute the maps of p -values in Figure 3.12, and the ΔR_2^* curves in Figure 3.11.

the baseline R_2^* by its effect on the signal S and the physiological component of the noise $\sigma_P = \lambda S$, it can be shown that ΔR_2^* is given by

$$\Delta R_2^* = -\frac{K \sqrt{\sigma_0^2 + \lambda^2 M_0^2 e^{-2TER_2^*}}}{M_0 e^{-TER_2^*} TE}, \quad (3.31)$$

which illustrates that the required ΔR_2^* depends on TE , the baseline R_2^* , the factor K and the relationship between the thermal noise σ_0 and physiological noise, σ_P , represented by λ . Figure 3.11 illustrates the signal change ΔS (bold lines) and R_2^* change ΔR_2^* (dashed lines) as a function of TE for $T_2^* = 66$ ms and $T_2^* = 33.2$ ms, corresponding to estimates in GM tissue at 3T and 7T respectively [283]. It can be observed that the demand in terms of ΔR_2^* to detect an activation is stricter at 7T than at 3T. Even though changes in CBV and CMRO₂ are physiologically constant across MR fields, the effect in cerebrovascular reactivity in terms of ΔR_2^* achieved at 7T is higher than at 3T so that higher sensitivity to the BOLD effect is feasible at higher MR fields [102].

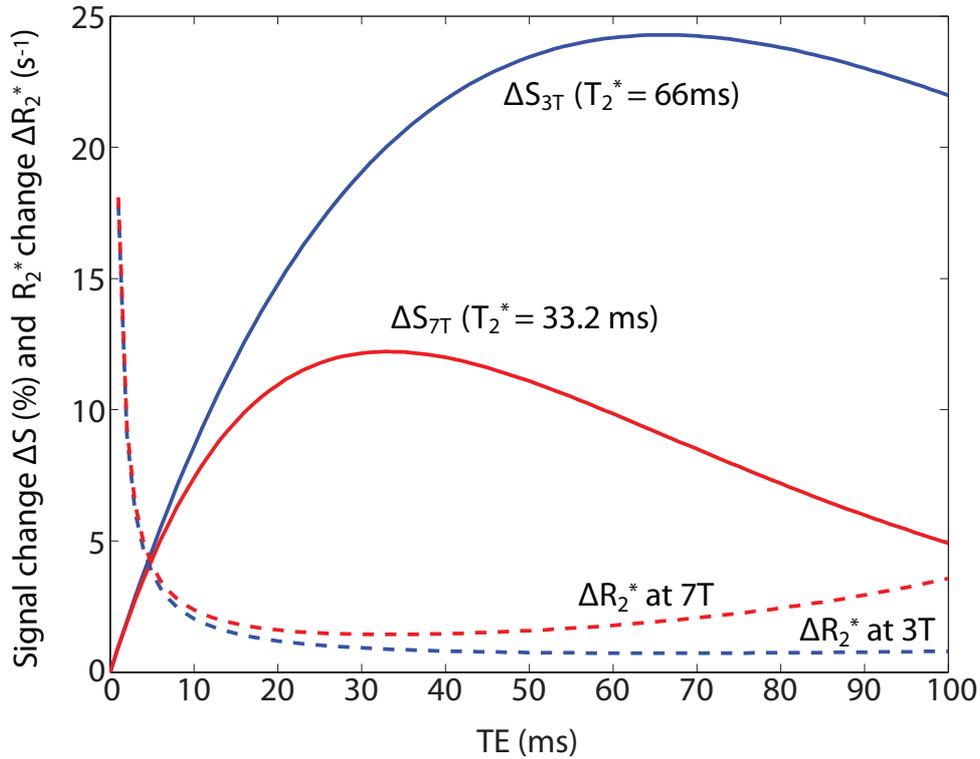


Figure 3.11: Signal change ΔS (bold lines) and R_2^* change ΔR_2^* (dashed lines) as a function of TE for $T_2^* = 66$ ms and $T_2^* = 33.2$ ms, corresponding to estimates in GM tissue at 3T and 7T respectively [283]. The curves for ΔR_2^* are computed for $K = 8.78$, and assuming $M_0 = 1$, $\lambda = 0.0112$, and $\sigma_0 = 2$.

In order to compensate for differences in BOLD sensitivity due to variations in baseline R_2^* , correction factors for the signal change and the noise can be computed as if all voxels were sampled at the optimal $T_2^* = TE$. Then, the t -statistics can be corrected based on Eqs. (3.28) and (3.30). Hereinafter, X_a and X_c will denote the values of the parameter X when the signal is measured at the actual echo time $t = TE$ and corrected $t = T_2^*$, respectively. Assuming $t = T_2^*$ the MR signal would be $S_c = M_0 e^{-1}$, whereas the signal acquired at $t = TE$ is $M_0 e^{-TE R_2^*}$. Therefore, we can define a signal correction factor f_s such that $f_s = S_c/S_a$:

$$f_s = e^{TE R_2^* - 1}. \quad (3.32)$$

Similarly, it can be shown that the correction factor for the BOLD signal change ΔS so that $f_{\Delta S} = \Delta S_c/\Delta S_a$ is also equal to f_s .

However, as noted in §3.4, the physiological term of the noise, σ_P , also scales with the MR signal change as a function of λ . Based on the noise model given in

Eq. (3.8), we can define a noise correction factor such that $f_\sigma = \sigma_c/\sigma_a$:

$$f_\sigma = \sqrt{\frac{\sigma_0^2 + \lambda^2 M_0^2 e^{-2}}{\sigma_0^2 + \lambda^2 M_0^2 e^{-2TER_2^*}}}. \quad (3.33)$$

Substituting for $SNR_0 = M_0 e^{-TER_2^*}/\sigma_0$, the noise correction factor can be written as

$$f_\sigma = \sqrt{\frac{1 + \lambda^2 SNR_0^2 e^{2TER_2^*-2}}{1 + \lambda^2 SNR_0^2}}. \quad (3.34)$$

Combining the factors for the signal and the noise, the t -statistics can be corrected for variations in baseline R_2^* by multiplying them with the following correction factor:

$$f_t = e^{TER_2^*-1} \sqrt{\frac{1 + \lambda^2 SNR_0^2}{1 + \lambda^2 SNR_0^2 e^{2TER_2^*-2}}}, \quad (3.35)$$

which depends on voxelwise estimates of the SNR_0 , λ , baseline R_2^* and TE .

3.6.2 Methods

In order to illustrate the concept of BOLD sensitivity maps, two datasets were acquired in one subject scanned during a hypercapnic challenge at 3T and 7T field strengths. We used this paradigm because it causes a uniform BOLD signal change across the whole cortex. The hypercapnic challenge consisted of an initial 3 min of baseline $P_{ET}CO_2$ (40 mm Hg) followed by a 2 min period of a target level of pressure of end-tidal CO_2 ($P_{ET}CO_2$), followed by 1 min of baseline with this cycle repeated for 5 cycles (targeted $P_{ET}CO_2$ time course) [102]. MR images were acquired on a Philips Achieva 3T MR scanner with a whole body volume transmit and 8-channel SENSE head receive coil, and a Philips Achieva 7T MR scanner with head volume transmit and 16-channel SENSE head receive coil. Nine slices of fMRI data were acquired with a double-echo, single-shot GE EPI sequence with $TE=16/81$ ms at 3T, and $TE=15/58$ ms at 7T, where the TE values were optimized for estimation of the baseline R_2^* . The total number of scans was 800 acquired with $TR=1.5$ s. The voxel size was $2 \times 2 \times 3$ mm³ and parallel imaging was used with a SENSE factor of 2. Tissue classification inversion-recovery echo-planar (IR-EPI) images were acquired with the same geometry at the null imaging points of GM, WM and CSF ($TI=600/900/1900$ ms at 3T, and $TI=600/1100/2200$ ms at 7T) [102].

Datasets were initially corrected for motion with AFNI (NIMH/NIH, USA).

Then, R_2^* time courses were calculated from the double-echo data using (3.7) and then corrected for low frequency drifts. Voxelwise estimate of R_2^* were computed as the average R_2^* during the middle baseline periods of the time course, creating R_2^* maps. The fMRI data acquired with the first echo time (TE=16 ms at 3T and TE=15 ms at 7T) were used to obtain a voxelwise estimate of $tSNR$. BOLD sensitivity maps of the ΔR_2^* required to detect significant activation were computed as $(tSNRTE)^{-1}$, based on Eq. (3.30) regardless of the constant factor K . Assuming a single trial model with stimulus duration 2 s, TR 2 s and 60 scans ($K = 8.78$, black dot in Figure 3.10), maps of the p -value resulting from a constant ΔR_2^* were also computed.

Furthermore, we computed t -statistics of the signal change generated by the hypercapnic challenge in the fMRI signal of the first echo time at 7T. For that, a GLM analysis was formed with two regressors: the time course of the targeted $P_{ET}CO_2$ time course and an all-ones vector to represent the mean. Correction factors for these statistics were computed based on Eq. (3.35) and were multiplied to the initial t -values. Since no data was acquired to obtain voxelwise estimates of λ [200, 201] and SNR_0 [190], we assumed constant values of $\lambda = 0.0112$ and $SNR_0 = 131.7$ for all voxels, as suggested in Triantafyllou et al. (2005) [343] for GM tissue with magnetic field strength of 7T and voxel resolution of $2 \times 2 \times 3$ mm³.

3.6.3 Results and discussion

Figure 3.12 shows the sensitivity maps of the required ΔR_2^* and p -values achieved with a constant ΔR_2^* of 5 s^{-1} for the single trial model (duration 2 s, TR 2 s, 60 scans) at 3T and 7T. The sensitivity maps and p -maps demonstrate that there exist significant variations in BOLD sensitivity across the cortex at both magnetic fields. At 3T, higher ΔR_2^* changes are required in frontal areas than in posterior areas, whereas higher ΔR_2^* is required at 7T in CSF areas, such as the lateral ventricles. Across magnetic fields, larger values of ΔR_2^* are required at 7T than at 3T as suggested in Figure 3.11. The maps also illustrate high BOLD sensitivity in WM tissue because the TE of the datasets was closer to WM than GM (T_2^* values at WM/GM: 53.2/66.0 ms at 3T; 26.8/33.2 ms at 7T [283]). Equivalently, the p -maps show that more significant p -values are achieved at 7T than at 3T, demonstrating the higher sensitivity available at higher fields. Furthermore, the maps also illustrate that set-

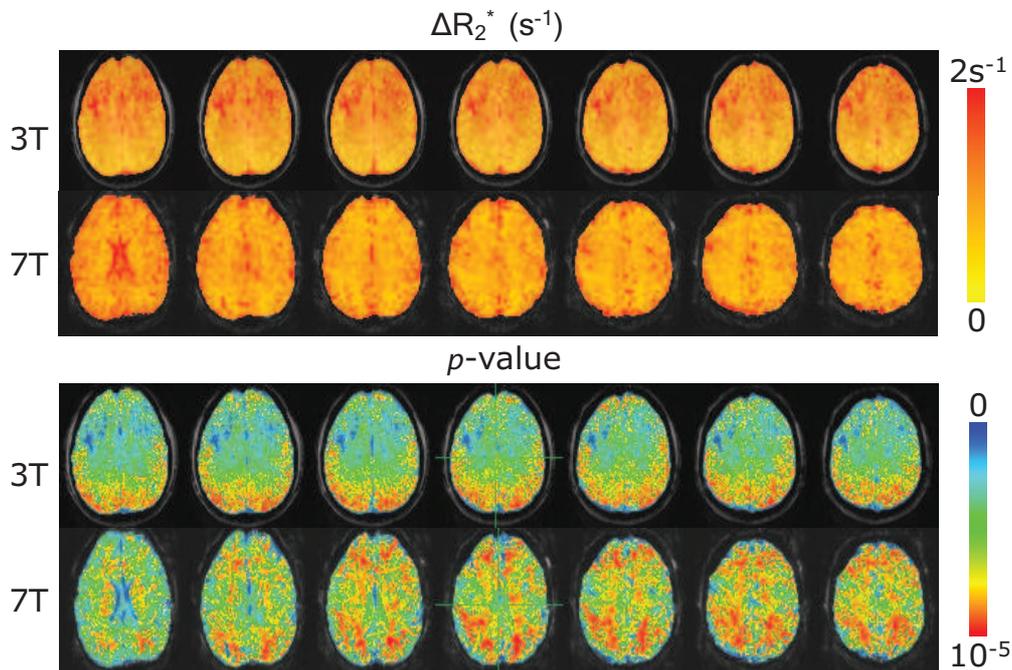


Figure 3.12: Sensitivity maps of the required ΔR_2^* (top) and p -values achieved with a constant ΔR_2^* of 5s^{-1} for the single trial model (duration 2 s, TR 2s, 60 scans) (bottom) at 3T and 7T.

ting a constant p -value threshold can obscure the detection of significant BOLD changes in areas with low sensitivity. For instance, the sensitivity maps at 3T illustrate that detecting BOLD effects in frontal cortical areas requires a larger change in blood flow and BOLD signal than in parietal or somatosensory regions. Consequently, by setting a statistical threshold which is appropriate to detect activations in the somatosensory areas (with relatively long T_2^*), simultaneous activations in the prefrontal lobes (with relatively short T_2^*) might not achieve enough statistical significance to be detected (i.e. becoming a false negative event). This nonuniformity in thresholding proves that sensitivity variations between the different regions of interest affect the statistical power of the final fMRI results [151]. Adequate shimming of the regions of interest [21, 288], field map correction [182, 183] or using parallel imaging techniques [91, 155, 292] can improve B_0 and B_1 homogeneities and diminish differences in BOLD sensitivity. Nevertheless, even if those strategies were employed or not employed at all, variations in BOLD sensitivity could remain and mapping them by means of BOLD sensitivity maps can facilitate any statistical interpretation of the results.

Figure 3.13 plots the correction factors for the dataset at 7T and its effect in

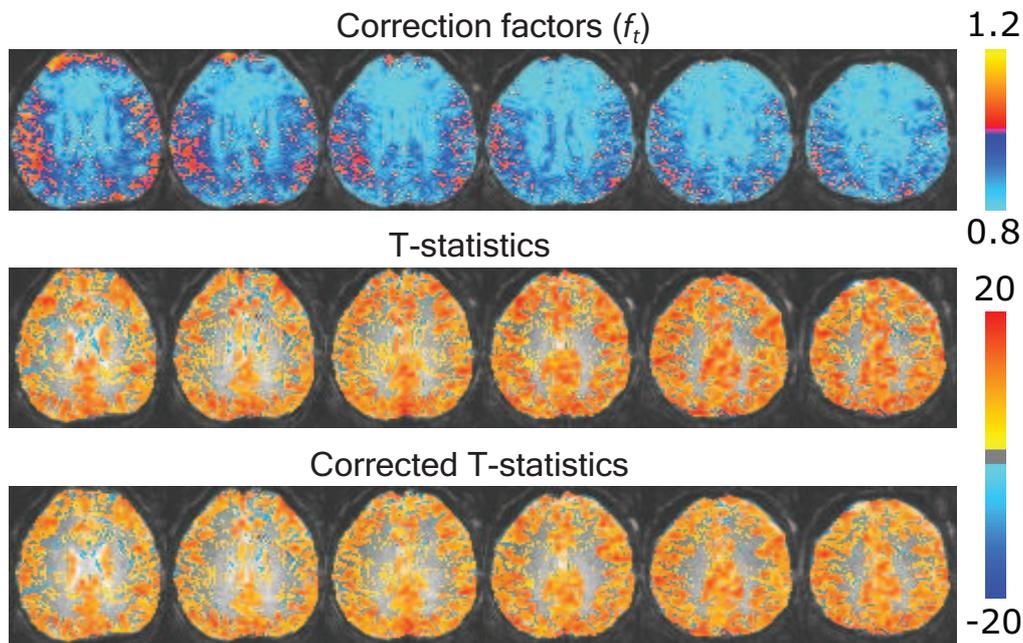


Figure 3.13: Correction factors (top), standard t -statistics (middle), and corrected t -statistics for the dataset of the hypercapnic challenge acquired at 7T.

the corresponding t -statistics of the hypercapnic challenge. Higher sensitivity to the hypercapnic paradigm is observed in GM than in WM (maps are thresholded at $p < 0.001$ FDR corrected), in agreement with [102]. The maps of correction factors illustrate that the t -statistics of frontal regions would scale down, whereas the statistics of posterior GM voxels would increase in case of an optimal T_2^* acquisition. Correcting the statistics has little effect in GM regions but does reduce the number of active voxels ($p > 0.001$, FDR corrected) in WM and CSF areas, in particular in those voxels showing negative t -values before correction. However, care must be taken in order to investigate the effect of a voxelwise estimation of SNR_0 and λ , along with more elaborate procedures for R_2^* estimation which compensate macroscopic magnetic field inhomogeneity effects [283]. In addition to correction of the statistics, maximal ratio combination of the signals acquired with a multiecho acquisition have been proven beneficial over single-echo acquisitions to reduce the T_2^* dependence of BOLD sensitivity [151, 291, 316].

3.7 Conclusion

After a brief review of functional imaging techniques, the chapter focused on describing the principles of BOLD fMRI. First, the physiological basis of the BOLD

effect and the temporal characterization of the associated haemodynamic response were described. In addition, the main concepts of BOLD sensitivity in gradient echo EPI were presented, along with a physiological model for the noise observed in BOLD fMRI. The chapter also drew attention to fMRI data analysis based on the GLM and ICA methods. Both techniques will be used in following chapters. Finally, we presented a novel approach to map BOLD sensitivity in terms of the necessary change in R_2^* to detect a statistically significant activation and introduced correction factors in order to account for the variability in R_2^* across brain regions.

The following chapters in this thesis will build on the concepts presented in the previous two chapters in order to describe three novel paradigm free mapping techniques to detect single trial event related BOLD responses without prior information of the stimulus timing.

Chapter 4

Paradigm Free Mapping

Throughout this thesis, novel fMRI methodologies are presented for mapping the brain's responses to single trial events or actions without prior information about the timing or locations of the events, Paradigm Free Mapping (PFM). The chapter begins with an introduction to the principles of PFM (§4.1), followed by a description of the theory in which the first PFM technique is based (§4.2). The technique is evaluated in a real fMRI experiment where subjects performed visually-cued and self-paced finger tapping tasks (§4.3 and §4.4). The chapter finishes with a brief discussion of the results and the theory of the technique presented in this chapter (§4.5). The ideas introduced in this chapter form the basis of the PFM extensions presented in Chapters 5 and 6.

4.1 Introduction to Paradigm Free Mapping

Paradigm Free Mapping (PFM) aims to detect and characterize the cortical responses to isolated events or actions with no prior information about the timing or locations of the events. The next sections give a brief review about the potential of both concepts in fMRI analysis.

4.1.1 The potential of single trial fMRI

Mapping the cortical response to a single mental event in both space and time is essential to improve our understanding of the basis of human perception and cognition. fMRI paradigms typically use event-related (ER) designs to characterize individual cognitive events and analyze the cortical response to isolated trials by measuring the hemodynamic response function (HRF) [51, 186]. The traditional

approach in ER-fMRI is to average across several single trial responses in order to improve estimation of the HRF parameters [30], but this assumes that the haemodynamic response and neuronal activity is constant across trials[54]. It is generally assumed that the improvement in the estimate of the HRF parameters is proportional to the square root of the number of trials averaged. However, there is intrinsic variability in the observed HRF between successive trials (inter-trial variability), across trials of different sessions (inter-session variability) and across subjects (inter-subject variability) [4, 103, 162]. Variability in the response may arise not only from physiological effects, motion, systematic confounds such as magnetic field inhomogeneities, but also from uncontrolled changes in the subject's attention, perception, cognition or performance [63, 356]. Although averaging responses across trials enhances the characterization of the common response to a task, it also avoids the characterization of particular differences in the subject's response to a task, which may be important for understanding finer aspects of mental processing, such as learning or adaptation [145, 154]. Parametric designs can be investigated to assess linear or nonlinear relationships between the haemodynamic response amplitude, neuronal activity and experimental parameters [62]. Inference about the timing of neuronal activity can be obtained by comparing the haemodynamic response across different conditions [30, 52]. In general terms, functional imaging at ultrahigh MR fields (7T) provides sufficient BOLD contrast to noise ratio to allow the detection and characterization of the response to true single trials in fMRI paradigms [286, 349].

Even though the idea of using single trial events in fMRI events was initially proposed by Buckner et al. in 1996 [54], the first studies of true single trial experiments were independently presented by Ritcher et al. in 1997, Zarahn et al. in 1997, and Menon et al. in 1998, and investigated the brain response during a delayed cued finger movement task [298], mental rotation [300], visual discrimination [381] and visuomotor tasks [249]. The technique, called time-resolved fMRI analysis, allows the investigation of the sequential orchestration of cortical responses across brain regions or mental chronometry of the brain [248, 249] by correlating the estimated hemodynamic response parameters, such as onset time or width, with behavioural parameters, such as reaction time. Since then, single-trial time-resolved fMRI analysis has been extensively used to study different cognitive processes with fMRI [37, 85, 123, 309, 381], and also in multimodal imaging with EEG-fMRI [93, 94]. Other popular techniques for single trial fMRI analysis rely on the deconvolution of the shape of the HRF [227], autoregressive models with exogenous inputs (ARX)

models [23], or fitting of a predetermined HRF function and parameters [199, 335]. Decisively, these single trial fMRI techniques assume prior knowledge of the stimulus or paradigm timing.

4.1.2 The potential of no timing information

In many studies an exact temporal model of the expected BOLD signal cannot be formulated or predicted in advance, such as clinical studies of hallucinations in schizophrenia [192], interictal discharges in epilepsy [17], or cognitive paradigms which involve switches in the subject's perception, such as visual priming [194] or binocular rivalry [341], where the subject's behavioural response is difficult to be accurately synchronized with the cortical response. Consequently, there is a need for approaches to analyze fMRI data which avoid the requirement to specify the onsets of cortical responses, and which allow the study and use of more unconstrained experimental paradigms, particularly for single-trial events. In addition, the study of task-unrelated brain activity may also provide novel information about transient, spontaneous cortical activations which characterize the basal brain state at rest [35, 124].

Model-free methods

If there is no information about when the activation occurs or the shape of the HRF, model-free techniques, such as clustering methods [25, 26, 117, 150], principal component analysis (PCA) [8, 24, 337], independent component analysis (ICA) [27, 65, 246] and temporal clustering analysis (TCA) [220, 255, 257] can be used. These methods have been proven efficient in identifying voxels showing consistent and common temporal or spatial patterns of activation (components), which later can be used as regressors in a General Linear Model analysis for statistical inference. Nevertheless, since the initial formulation does not rely on any model, the components are susceptible to noise and scanner artifacts. Furthermore, the results of the analysis could greatly vary depending on the number of components or clusters estimated from the data, which are defined a-priori or estimated from the data [163, 193], or the initialization of the algorithm in the case of ICA [168]. As for TCA methods, the technique couples together voxels based on the timing of its first maximum or activation, and so this approach would not be useful if the initially coupled regions were activated in response to two different stimuli [257].

Regarding the application of data-driven techniques for single-trial paradigms,

ICA-based methods have been demonstrated to be very promising to study the variability of BOLD responses [103], successful for predicting human errors in combination with a finite impulse response deconvolution [109], and able to display the mental chronometry of auditory-cued and visual imagery task [326]. As for clustering techniques, self-organizing mapping [263] or fuzzy clustering [299] obtained comparable performance as time-resolved fMRI in a sequential finger movement task studied by Ritcher et al. (1997) [298].

Model-based methods for PFM

Contrary to popular belief, model-based techniques have also been developed to study the brain's response with no timing constraints. To begin with, although the traditional GLM analysis mainly focuses on the estimation of the amplitude of the activations, inference about the delay or latency of the response can be still obtained by using the temporal and dispersion derivative of the HRF to account for shifts of one or two seconds [67, 130, 167, 213, 214, 374]. With similar reasoning, Ciuciu et al. (2003) proposed the Bayesian deconvolution of an asynchronous HRF model in which the image acquisition rate is jittered with the stimuli, allowing for small timing shifts in the response relative to the stimulus [78]. Note, however, partial or total a-priori information about the paradigm timing is necessary for the definition of the model regressors.

Probabilistic frameworks based on Hidden Markov models have been presented in Faisan et al. (2007) [118] and Hutchinson et al. (2009) [176]. Hidden Markov multiple event sequence modelling (HMMESM) [118] enables spatio-temporal mapping of the response by first searching the voxel time series for spikes, which are then incorporated into a Hidden Markov chain of candidate onsets associated with a hemodynamic response. However, no statistical inference is given on the activation onsets. Alternatively, Hutchinson et al. (2009) [176] proposed to disentangle a sequence of multiple hidden processes by estimating their corresponding spatio-temporal responses, in addition to the probability distribution of the lag between stimulus and activation of each process. Nevertheless, it is recommended that the stimulus onsets should be specified in order to diminish the computational complexity of the technique, with similar recommendations in Ciuciu et al. (2003) [78].

From a deconvolution point of view, Riera et al. (2004) proposed blindly estimating the dynamics of the physiological components driving the BOLD responses under an extended Ballon-model and simultaneously reconstructing the underlying

activations. However, the practical application of this approach for single-trial fMRI is challenging due to the large amount of parameters to be estimated. Based on the L_1 -norm deconvolution, an HRF-shaped wavelet analysis, termed 'activelets' [191], has shown promise in detecting simulated single-trial event-related responses with no prior timing information.

Finally, change point theory methods enable the characterization of prolonged single-epoch paradigms without timing information by modelling the voxel time series as the mixture of two Gaussians distributions, one for the baseline state and another for the activations [215, 303]. Interestingly, Lindquist et al. (2007) introduced a methodology for single-subject analysis where the data was processed with an Exponentially Weighted Moving Average (EWMA) filter and then significant changes in the filtered time series were detected by a Hotelling T^2 -statistic against a baseline state [215]. The PFM method presented in this chapter follows this last idea in order to detect significant activations (see §4.5).

4.2 Theory

All the PFM techniques presented in this thesis involve two-stages of analysis. The first stage is based on the voxel-wise deconvolution of the HRF from the fMRI time series assuming linearity in the BOLD response [48, 143]. In the second stage the statistical significance of the deconvolved responses is determined. Consequently, PFM must be understood as a model-based (model of HRF) approach which enables the study of brain activity associated with single trial events in a simple and exploratory manner, without any knowledge of the paradigm timing.

Chronologically, the technique described in this chapter was the first approach based on the PFM principles, and consequently it will be referred to as PFM in the following. The technique computes the deconvolution of the HRF by means of the ridge regression estimator [170], assuming that the noise is an autoregressive (AR) stochastic process [333]. Statistical inference is performed by t -statistics which assess the significance of signal deviations from a baseline state for each voxel and time point [215]. The t -statistics are finally corrected for spatial dependence with the Benjamini-Hochberg (BH) procedure to control the false discovery rate (FDR) [32].

4.2.1 Signal model

In the absence of noise and assuming a linear, time-invariant model [48], the continuous BOLD signal, $y(t)$, in a voxel can be modelled as the convolution of an input neuronal-related signal, $s(t)$, and the haemodynamic response function (HRF), $h(t)$,

$$y(t) = h(t) * s(t) = \int_{\tau} h(\tau) s(t - \tau) d\tau. \quad (4.1)$$

The continuous signal is then sampled during the MRI acquisition at successive TRs ($t = n\text{TR}$) such that we can write a discrete model for the BOLD signal as

$$y_n = h_n * s_n = \sum_{l=0}^{L-1} h_l s_{n-l}, \quad (4.2)$$

for time points $n = 0, \dots, N - 1$, where N is the number of time points, L is the discrete length of the HRF, and y_n denotes the sample of $y(t)$ at time $n\text{TR}$. Throughout this thesis, it is assumed that the coefficients of \mathbf{h} are normalized such that $\|\mathbf{h}\|_2^2 = 1$. This discrete time model can be rewritten in matrix formulation as

$$\mathbf{y} = \mathbf{H}\mathbf{s}, \quad (4.3)$$

i.e.

$$\begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_{N-1} \end{bmatrix} = \begin{bmatrix} h_0 & 0 & \dots & 0 \\ \vdots & h_0 & \ddots & 0 \\ h_{L-1} & \vdots & \ddots & 0 \\ 0 & h_{L-1} & \dots & h_0 \end{bmatrix} \begin{bmatrix} s_0 \\ s_1 \\ \vdots \\ s_{N-1} \end{bmatrix}, \quad (4.4)$$

where \mathbf{y} and \mathbf{s} are column vectors of length N denoting the measured fMRI and neural-related signals, respectively. The matrix \mathbf{H} is the convolution (Toeplitz) matrix of size $N \times N$ defined from the HRF.

In practice, the fMRI signal also includes confounds of systematic or physiological origin, such as scanner drifts or cardiac and respiratory fluctuations, which will be incorporated into the signal model as an additive noise term $\epsilon(t)$ (denoted as ϵ in discrete time). Thereby, the signal model is given by

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \epsilon. \quad (4.5)$$

In this chapter, the noise is assumed to be a stationary, stochastic process following a Multivariate Normal distribution with zero mean and temporal correlation

matrix, $E[\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T] = \sigma^2\boldsymbol{\Sigma}$, i.e. $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \sigma^2\boldsymbol{\Sigma})$, where the variance σ^2 and autocorrelation matrix, $\boldsymbol{\Sigma}$, are assumed to vary across voxels, i.e. voxel-dependent, in order to account variations in the noise statistical characteristics observed across regions and tissues [369, 373].

4.2.2 Deconvolution with L_2 -norm regularization: From Least Squares to Ridge Regression

The analysis with PFM is based on the assumption that the underlying neuronal signal \mathbf{s} provides more accurate information about the dynamics of brain activation than the hemodynamic response observed in BOLD fMRI [142]. Therefore, the goal is to estimate \mathbf{s} by means of the deconvolution of the HRF.

Maximum likelihood or least squares deconvolution

Since we assume that the noise follows a Gaussian distribution, the probability distribution of the signal is $p(\mathbf{y}|\mathbf{s}, \sigma, \boldsymbol{\Sigma}) = \mathcal{N}(\mathbf{H}\mathbf{s}, \sigma^2\boldsymbol{\Sigma})$. An estimate of \mathbf{s} can be obtained by maximizing the likelihood or probability of the voxel time series, which is called the Maximum Likelihood (ML) estimator [166]:

$$\hat{\mathbf{s}}_{ML} = \arg \max_{\mathbf{s}} p(\mathbf{y}|\mathbf{s}, \sigma, \boldsymbol{\Sigma}) \quad (4.6)$$

where

$$p(\mathbf{y}|\mathbf{s}, \sigma, \boldsymbol{\Sigma}) = (2\pi\sigma^2)^{-N/2} |\boldsymbol{\Sigma}|^{-1/2} \exp\left(-\frac{1}{2\sigma^2} (\mathbf{y} - \mathbf{H}\mathbf{s})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{H}\mathbf{s})\right). \quad (4.7)$$

It can be seen that the ML estimator (4.6) is equivalent to minimizing the negative logarithm of the likelihood such that

$$\hat{\mathbf{s}}_{ML} = \arg \min_{\mathbf{s}} (\mathbf{y} - \mathbf{H}\mathbf{s})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{H}\mathbf{s}). \quad (4.8)$$

Therefore, maximizing the likelihood for a Gaussian distribution is equivalent to the Generalized Least Squares (GLS) estimator which minimizes the residual sum of squares for a general covariance matrix $\boldsymbol{\Sigma}$. It can be shown that the GLS solution is given by

$$\hat{\mathbf{s}}_{GLS} = (\mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H})^{-1} \mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{y}, \quad (4.9)$$

and if the noise observations are uncorrelated, i.e. $\boldsymbol{\Sigma} = \mathbf{I}$, being \mathbf{I} the $N \times N$ identity

matrix, the GLS estimator reduces to the ordinary least squares (OLS) estimator, or simply least squares (LS):

$$\hat{\mathbf{s}}_{OLS} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{y}. \quad (4.10)$$

Nevertheless, there are some cases where the GLS estimator (4.9) may not be the optimal one. First, observe that in the signal model (4.5) the number of regressors in \mathbf{H} is equal to the number of observations N . Consequently, the ML estimate has a high probability of exactly fitting the observations so that all the coefficients of $\hat{\mathbf{s}}_{ML}$ are nonzero. As a result, the ML estimates yield an accurate fit (low bias) but are very inefficient in terms of prediction error (large variance). The large variance is especially noticeable at high frequencies because the estimates at low frequencies are controlled by the smoothness of the HRF [142].

The estimation can be improved by incorporating some kind of penalization or regularization to the least squares function (4.8) so that simpler models are favoured. In linear models such as (4.5), reducing the model complexity translates to shrinking some coefficients of \mathbf{s} to zero. By doing that, the variance of the fitted values is reduced at the expense of allowing some bias in the estimation of \mathbf{s} . Consequently, these type of estimators are often called shrinkage, regularized or penalized estimators [166].

The regularization term is generally established in terms of the L_p -norm of the coefficients:

$$\|\mathbf{s}\|_p = \left(\sum_i |s_i|^p \right)^{1/p}. \quad (4.11)$$

For instance, L_2 for $p = 2$ is the Euclidean norm of a vector. Figure 4.1 plots the countours of the L_p -norm for different values of p for a two-dimensional vector. In practice, L_p -norms with $0 < p \leq 2$ are commonly employed and it is well-known that a selection of $0 < p \leq 1$ is appropriate when one expects that most of the coefficients are equal or very close to zero, i.e. promoting that \mathbf{s} is a sparse vector with few coefficients significantly different from zero. In fact, the L_0 -norm is a measure of sparsity in the strictest sense since it is defined as the number of nonzeros in \mathbf{s} [50]. In this chapter, the Euclidean norm or L_2 -norm is used to regularize the estimate of \mathbf{s} , whereas L_1 -norm estimators are investigated in Chapters 5 and 6.

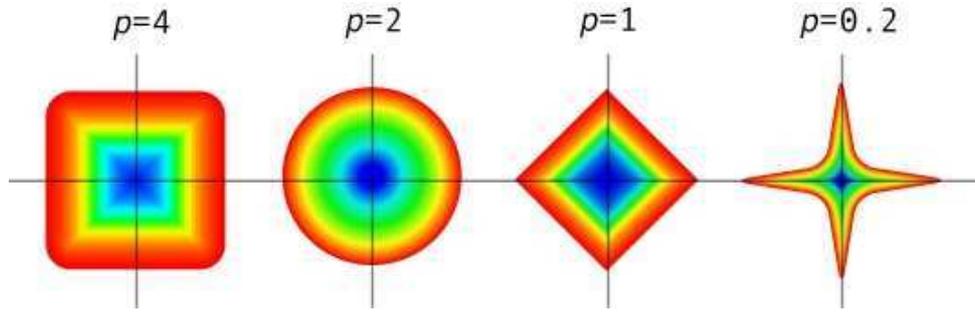


Figure 4.1: Contours of the L_p -norm for different values of p . Each colour indicates a constant maximum value of the corresponding L_p -norm. These figures can also be understood as prior distributions of the coefficients for the two-dimensional case. Shown above are a very uniform ($p=4$), Gaussian ($p=2$), Laplacian ($p=1$) and a very sparse distribution ($p=0.2$).

Ridge Regression deconvolution

The shrinkage method of ridge regression (RR) [170] estimates the signal \mathbf{s} by imposing an L_2 -norm penalty on the coefficients. The generalized RR estimate is computed as

$$\hat{\mathbf{s}}_{RR} = \arg \min_{\mathbf{s}} (\mathbf{y} - \mathbf{H}\mathbf{s})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{H}\mathbf{s}) + \lambda \mathbf{s}^T \mathbf{s}, \quad (4.12)$$

where λ is a nonnegative regularization parameter which controls the tradeoff between minimizing the least squares term and the penalization term. The larger the value of λ , the greater the amount of shrinkage of the coefficients. It can be shown that (4.12) is the minimization of the Lagrange function:

$$\hat{\mathbf{s}}_{RR} = \arg \min_{\mathbf{s}} (\mathbf{y} - \mathbf{H}\mathbf{s})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{H}\mathbf{s}) \quad \text{subject to} \quad \mathbf{s}^T \mathbf{s} \leq \gamma \quad (4.13)$$

This equivalent formulation of the problem is useful to understand how the RR estimator operates and how it is useful in the proposed model. Notice that, by definition of the convolution (4.3), the columns of \mathbf{H} are non-orthogonal and the matrix $\mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H}$ may be singular or non-invertible. Therefore, if the GLS estimator were to be used, the estimates become poorly determined and exhibit high variance. For instance, a very large positive coefficient could be cancelled out by another very large negative coefficient from a correlated regressor. By imposing a constraint in the L_2 -norm of the coefficients as shown in (4.13), the variability of the coefficients and the complexity of the model are controlled [166].

It can be shown after some algebra that the RR estimate from (4.12) can be

written as

$$\hat{\mathbf{s}}_{RR} = (\mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H} + \lambda \mathbf{I})^{-1} \mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{y}, \quad (4.14)$$

where the RR solution becomes equal to the GLS solution (4.9) when $\lambda = 0$.

Therefore, applying a L_2 -norm regularization translates to adding a diagonal constant λ which prevents the matrix $(\mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H} + \lambda \mathbf{I})$ from being singular even if the matrix $\mathbf{H}^T \mathbf{H}$ (or $\mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H}$) is not full-rank. This was the original rationale for the RR estimator when first proposed by Hoerl and Kennard in 1970 [170]. Furthermore, diagonal loading has been shown to enhance the robustness of linear estimators against model mismatches [226, 317]. In addition, the RR solution can be interpreted from a Bayesian perspective as the maximum a-posteriori estimation (MAP) of \mathbf{s} when the prior distribution of \mathbf{s} is Gaussian with zero-mean and uncorrelated coefficients, i.e. $\mathbf{s} \sim N(0, \beta^2 \mathbf{I})$, as shown in Appendix A.1.

Selection of the regularization parameter

The choice of λ is important to ensure that the RR algorithm computes estimates with low bias and low variance. In this study, following the Bayesian arguments described in Appendix A.1, the regularization parameter is set to [206],

$$\lambda = \frac{N \hat{\sigma}_{GLS}^2}{\hat{\mathbf{s}}_{GLS}^T \mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H} \hat{\mathbf{s}}_{GLS}}, \quad (4.15)$$

where $\hat{\sigma}_{GLS}^2 = \|\mathbf{y} - \mathbf{H} \hat{\mathbf{s}}_{GLS}\|_2^2 / (N - \text{rank}(\mathbf{H}))$ is the variance of the residuals after first fitting the GLS estimate with $\lambda = 0$ (note that the rank of \mathbf{H} is $N - 1$ and therefore the denominator is equal to 1). Hence, the regularization parameter is set to the quotient between the GLS estimates of the variance of the noise and the coefficients.

Alternative selection criteria for the regularization parameter could be considered such as L-curve [164], cross-validation methods [166] or using empirical bayes procedures to estimate the variance of the noise and the coefficients [142]. Nevertheless, in all these approaches the computation of λ is done iteratively which considerably increases the computational cost of the algorithm with small improvement over selecting λ with (4.15) in terms of robustness against model mismatches [317].

Statistical characterization of the noise

The signal model used in PFM considers that the noise exhibits temporal correlation which is modelled with the autocorrelation matrix Σ . The fMRI signal can be significantly affected by low-frequency fluctuations of systematic origin, such as drifts due to scanner instabilities [321], and physiological origin, such as cardiac and respiratory fluctuations [38, 71, 144, 320]. Modelling the temporal correlation of the noise is important in any fMRI analysis technique since conclusions based on test statistics or p -values could be largely biased when the noise is correlated but assumed to be an uncorrelated process [322].

Numerous strategies have been proposed to deal with temporal correlations in fMRI. A simple approach is to remove low-frequency fluctuations during preprocessing by regressing out a set of basis functions, such as sinusoidal signals, splines, wavelets [336]. This helps to reduce the most relevant low-frequency components. Even so, it is still recommendable to estimate the temporal correlation of the residual noise. For that, one can estimate the autocorrelation matrix and then 'prewhiten' the fMRI time series so that the noise can now be considered uncorrelated or 'white' [239, 369, 373]. Alternatively, one can shape or impose a predetermined autocorrelation via temporal filtering or smoothing [132, 372]. At first sight, smoothing or colouring approaches might be adequate in situations where it might be difficult to accurately know or estimate the intrinsic autocorrelation of the noise and so shaping the autocorrelation function can reduce the bias and variance of the estimates [135]. Nevertheless, in practice, the autocorrelation imposed by colouring does not completely shape the intrinsic autocorrelation present in the data [369]. In addition, it has been shown that the efficiency of the estimator with a prewhitening approach is higher than with a colouring method, especially for event-related designs [369]. As shown in Appendix A.2, the GLS estimator can be interpreted as a prewhitened OLS estimator.

In this study, the matrix Σ is estimated using the data acquired during a baseline or pre-stimulus period, which is assumed to have no signals or events of interest (see Figure 4.5). Furthermore, a spatial neighbourhood of K voxels is considered in order to reduce the variance of the estimate and incorporate spatial information [140, 371, 373]. First, the raw autocorrelation coefficients for a given voxel are

computed over a local neighbourhood

$$r(i) = \frac{1}{K} \sum_{k=1}^K \frac{1}{B-i} \sum_{n=1}^{B-i} y_k(n) y_k(n+i), \quad i = 0, \dots, B-1, \quad (4.16)$$

where $y_k(i)$ is the fMRI observation at time i of the baselines for the k^{th} voxel of the neighbourhood. Then, the matrix Σ is defined as the Toeplitz matrix of the autocorrelation coefficients $r(i)$.

Despite spatial averaging, using the raw autocorrelation coefficients directly is not advisable since those estimates will have large variance, especially at higher time lags, making the estimate of Σ and in turn the estimate of \mathbf{s} considerably biased. One solution is to smooth the raw correlation coefficients with a single or multitapering approach or characterize them with a parametric model (see for instance [46, 49, 333] for textbooks on the topic, and [239, 322, 369] for fMRI time series analysis). In this study, a parametric approach is implemented where the noise is modelled as an autoregressive process of variable order p (AR(p)),

$$\epsilon(n) + a_1 \epsilon(n-1) + a_2 \epsilon(n-2) + \dots + a_p \epsilon(n-p) = w(n) \quad (4.17)$$

where a_i ($i = 1, \dots, p$) are the AR model parameters and $w(n)$ is a white noise process with variance σ_p^2 . AR models are preferred to more complex models, such as autoregressive-moving average (ARMA) models, because any ARMA model can be closely approximated with an AR model with higher number of parameters and AR estimation algorithms exhibit higher accuracy and computational robustness than those for ARMA model estimation [49, 333]. Besides, AR models have been proven sufficient to model temporal correlations in fMRI when prior detrending is carried out to remove the most relevant low-frequency trends [280].

The Levinson-Durbin (LD) algorithm [333] is used here to estimate the AR model parameters for each candidate order p (a_i , $i = 1, \dots, p$) and σ_p^2 (see Appendix A.3). Once the AR parameters for each candidate order have been found, the optimal order, \hat{p} , of the AR model is selected based on statistical model selection criteria [333]. Under the assumption of Gaussian noise, model selection criteria find the optimal order which minimizes the variance of the residuals after fitting the AR model for order p , given by $\hat{\sigma}_p^2$, plus a penalty function, $K(p, B)$, which depends on the number of parameters to fit the model, p , and the number of observations, B .

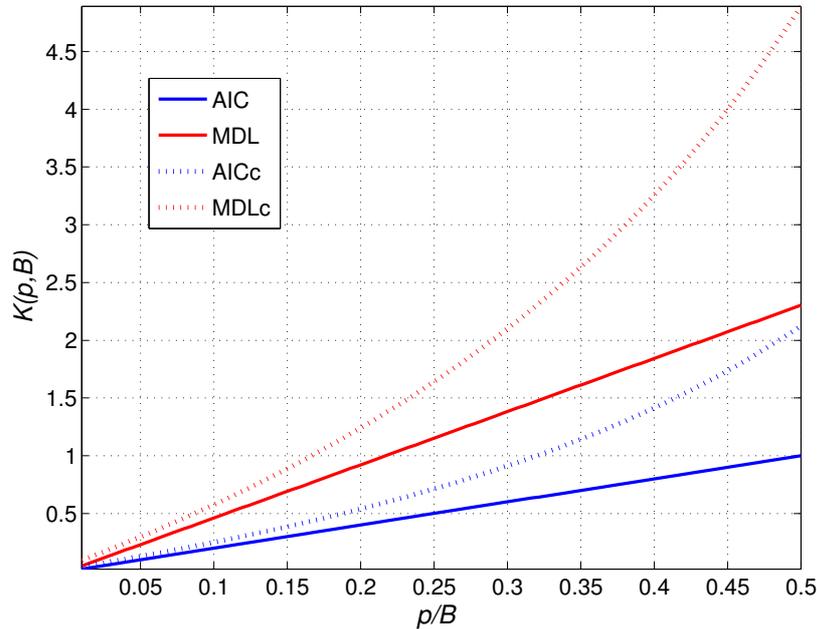


Figure 4.2: Penalization curves for the AIC, MDL, AICc and MDLc model selection criteria as a function of the ratio between the number of parameters of the model, p , and the number of observations, B . Finite sample criteria (MDLc and AICc, dotted lines) apply larger penalizations than the corresponding asymptotic criteria, which assume that $B \rightarrow \infty$. Besides, MDL-based criteria apply larger penalizations than AIC-based criteria and consequently tend to choose less complex models.

	AIC	AICc	MDL	MDLc
$K(B, p)$	$\frac{2p}{B}$	$\frac{2(p+1)}{N-p-2}$	$\frac{\text{Ln}Bp}{B}$	$\frac{\text{Ln}N(p+1)}{N-p-2}$

Table 4.1: Penalization factors for AIC, AICc, MDL and MDLc model selection criteria.

Mathematically, the optimal order \hat{p} is selected as

$$\hat{p} = \arg \min_p \quad \text{Ln}(\hat{\sigma}_p^2) + K(p, B). \quad (4.18)$$

In this study, asymptotic and finite-sample versions of the Akaike Information Criterion (AIC) [5, 175] and the Bayesian Information Criterion (BIC) [315, 301], which is equivalent to the Minimum Description Length (MDL) [302], were investigated. The terms $K(p, B)$ for these criteria are given in Table 4.1. Figure 4.2 plots the penalization curves as function of p/B . It can be seen that the MDL-based criteria penalize further those models with larger number of parameters than AIC-based, and so MDL-based criteria are more efficient in controlling model complexity and avoiding problems of overfitting [49].

The performance of the different model selection criteria in Table 4.1 was evalu-

ated based on the spatial maps of the optimal model orders and AR model parameters, and histograms of the AR model orders for each dataset (see §4.3 and §4.4). Based on these results, the MDLc criterion showed larger spatial coherence in the AR model parameters and consequently the optimal order \hat{p} for the deconvolution was selected as

$$\hat{p} = \arg \min_p \quad \text{Ln}(\hat{\sigma}_p^2) + \frac{\text{Ln}B(p+1)}{B-p-2}. \quad (4.19)$$

Finally, the estimate of the autocorrelation coefficients $r(i)$, $i = 0, \dots, N-1$, which define Σ , are calculated from the AR coefficients and the reflection coefficients corresponding to the optimal model order \hat{p} (see Appendix A.4).

4.2.3 Statistical Inference and Thresholding

4.2.4 Test-statistics

To assess the statistical significance of the activations, a t -statistic time course is defined from the RR estimates given in (4.14), comparing the signal at each time point to the mean of the baseline period for each voxel [215]. Significant deviation from the mean estimated during the baselines would cause the null hypothesis of non-activation to be rejected for a given voxel and time point. Hereinafter, we do not consider the subscript RR and the superscript k to simplify the notation. For each time point i after the baseline period, the t -statistic time series is computed over the L nearest neighbouring voxels as

$$t(i) = \frac{\hat{s}_L(i) - \hat{\mu}_L}{\hat{\sigma}_L \sqrt{1 + \frac{1}{B}}} \quad i > B, \quad (4.20)$$

where $\hat{s}_L(i)$ is the spatial mean of the ridge regression estimates at voxel k at time i , and $\hat{\mu}_L$ and $\hat{\sigma}_L$ are the mean and standard deviation of \hat{s}_L of the baseline period ($1 \leq i \leq B$) taking into account the correlation of deconvolved time series, all calculated over the neighbouring L voxels. Note that the t -statistics are not independent in time. Each temporal t -statistic has a marginal Student's t -distribution with $B-1$ degrees of freedom, assuming the observations in the baseline are independent of that at time point i ($i > B$) (see Appendix A.5 for full details). Finally, p -values are computed from the t -statistics at each voxel and at each time point after the baseline.

Multiple hypothesis correction

The total number of hypothesis tests is equal to the number of voxels being tested multiplied by the number of time points of the time series (excluding the baseline time points). This number is much greater than conventional statistical parametric approaches where a single statistic is computed per voxel. Due to the temporal and spatial dependence of the t -statistics and the large number of tests, using a Bonferroni correction would be very conservative, resulting in very low sensitivity in detecting activations. Therefore, in this study we propose to correct p -values in space for each time point individually by applying the traditional Benjamini-Hochberg (BH) procedure for false discovery rate (FDR) control [32].

The FDR procedure controls the expectation or mean of the false positive proportion, i.e.,

$$FDR = E \left[\frac{Q}{V} \right], \quad (4.21)$$

where Q is the number of false positives within the V tested hypotheses.

In order to compute the threshold to achieve a given FDR, the BH FDR procedure proceeds as follows. Let V be the total number of tested voxels and α be the desired rate of false positives tolerated in our test ($0 < \alpha < 1$) for each time point. First, order the p -values in increasing order such that [32]

$$p_{(1)} \leq p_{(2)} \leq \dots \leq p_{(V)}. \quad (4.22)$$

Next, we compute the FDR threshold, p_{FDR} , as the largest of the ordered $p_{(i)}$, i.e. $p_{FDR} = \arg \max_i \{p_{(i)}\}$, for which

$$p_{(i)} \leq \alpha \frac{i}{V}. \quad (4.23)$$

Finally, reject the null hypothesis of non-activation in the voxels with $p_{(i)} \leq p_{FDR}$, or equivalently, threshold the map of t -statistics at the t -value corresponding to p_{FDR} (see Figure 4.3). Another possible interpretation of the FDR procedure is terms of the corrected p -values, known as Q -values, which for a given $p_{(i)}$ can be computed as [380]

$$Q(p_{(i)}) = \min_{i \leq k} \left[p_{(k)} \frac{V}{k} \right]. \quad (4.24)$$

The null-hypothesis of non-activation is rejected in those voxels with FDR-corrected Q -values $< \alpha$ [380], as illustrated in Figure 4.3. Since PFM performs

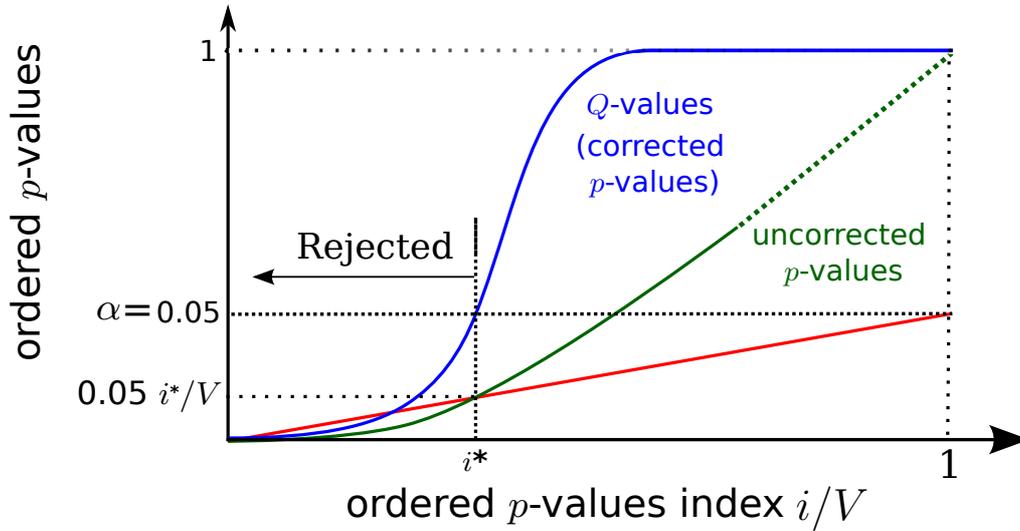


Figure 4.3: False discovery rate (FDR) procedure for $\alpha=0.05$. The uncorrected p -values are ordered increasingly (green line) and plotted as a function of i/V . The null-hypothesis is rejected for all those voxels with $p_{(i)} < \alpha i/V$ (red line). This is equivalent to computing the corrected p -values, Q -values, (blue line) and reject those voxels with $Q(p_{(i)}) < \alpha$.

the FDR correction for each time point independently, a time series is computed with the t -threshold for each time point given a FDR rate, e.g. $\alpha=0.05$.

4.3 Methods

Paradigm Free Mapping was evaluated with real data where subjects performed a visuo-motor paradigm with single trial events. Experiments were conducted with approval from the University of Nottingham ethics committee and informed consent was obtained from all participants.

4.3.1 Experimental paradigms

Six volunteers (5 male and 1 female, age 24-32 years, 5 right-handed and 1 left-handed) participated in the experiment. Functional MRI scans were acquired with BOLD acquisition at TRs of 2 s and 0.4 s. Figure 4.4 plots a schematic of the events and corresponding timing of the paradigm. The same paradigm was employed for all subjects at both TRs. Each scan started with an initial baseline (pre-stimulus) period of 140 s. This resulted in $B = 70$ baseline time points for TR 2 s and $B = 345$ baseline time points for TR 0.4 s. After this initial baseline, the subjects were visually cued to perform two trials of finger tapping (VCT) at 140 s and 180 s, each finger tapping trial being 4 s in duration. After 384 s, a message 'TAP at

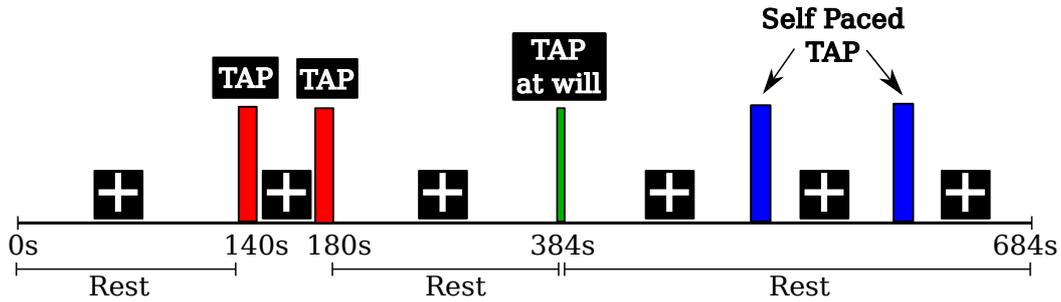


Figure 4.4: Visuomotor experimental paradigm. After the initial baseline period of 140 s, the subjects were visually cued to perform two trials of finger tapping at 140 s and 180 s, each finger tapping trial being 4 s in duration. After the message 'TAP at will' at time 384 s, the subjects were asked to carry out two trials of self-paced tapping (SPT) of similar duration to the VCT task. Throughout the periods of rest, subjects were instructed to fixate on a cross projected on the screen. Note that during the second part of the experiment, the fixation cross was continuously projected on the screen.

will' was projected onto the screen indicating the start of the second period during which subjects were asked to carry out freely two trials of self-paced tapping (SPT) of similar duration to the VCT task. Throughout the scan duration, subjects were instructed to fixate on a cross projected onto the screen when no other instruction was presented. Subjects were instructed about the paradigm prior to the scanning session. The visual instructions were projected from an LCD projector onto a screen located inside the scanner room, which subjects viewed through prism glasses with angle mirrors. The total scan duration was 684 s (342 scans for TR 2 s, after 5 scans discarded to achieve steady-state magnetization, and 1710 scans for TR 0.4 s, after 30 scans discarded).

4.3.2 MR data acquisition

MR images were acquired on a Philips 7T Achieva scanner (Best, Netherlands) using a 16-channel SENSE head coil (Nova Medical Inc., MA). Subjects' heads were secured in place using foam pads to minimize head motion. Cardiac and respiratory data were recorded throughout the fMRI acquisition using a respiratory belt and a pulse oximeter to allow physiological noise correction of the data. Datasets were acquired using single-shot, gradient echo EPI (TE: 30 ms, in-plane resolution: $2 \times 2 \text{ mm}^2$, slice thickness: 2 mm, SENSE factor: 1.5) with 20 slices for TR 2 s and 6 slices for TR 0.4 s. The flip angle was set to 80° or 40° , which approximated the Ernst angle for each TR respectively. At TR 2 s the imaging slices were positioned at approximately $+15^\circ$ to the cantho-meatal line [87] above the corpus callosum, to cover the supplementary motor area (SMA), premotor (PM)

and primary motor (M1) cortices, parietal cortex and the calcarine fissure. At TR 0.4 s a more tilted angle and 1 mm gap between slices was employed to cover approximately the same brain regions in fewer slices. Following the fMRI acquisition, high-resolution 1mm-isotropic resolution 3D anatomical T1-weighted (MPRAGE sequence) and T2*-weighted (spoiled-FLASH sequence) images were acquired to aid the localization of activations in functional and anatomical brain regions.

4.3.3 Surface electromyography (EMG) recording and analysis

Surface electromyography (EMG) was used to monitor the muscle activity in both hands during the fMRI experiment. Active electrode pairs were positioned on the left extensor (LE), right extensor (RE) and right flexor (RF) digitorum with paired electrode wires twisted to minimize the differential effect of the magnetic field on the EMG leads [352]. Ground and reference electrodes were positioned on bony parts of the wrist. Although subjects were instructed to perform finger tapping with their dominant hand (five right-hand subjects and one left-hand subject), EMG signals from both hands were recorded. EMG recording was performed using a Brain Vision recorder and a MR compatible BrainAmp amplifier (Brain Products, Munich, Germany). The amplifier was placed at the foot of the scanner bed on foam to minimize the effects of vibration. EMG signals were transmitted via an optical cable and recorded on a PC outside the scanner room. Data were sampled at 5000 Hz/channel and EMG analysis performed using Brain Vision Analyzer software. The quality of the EMG signals was checked visually prior to scanning and during the fMRI acquisition. EMG data were corrected for scanner artifacts using the average artefact subtraction method [6]. The differential signal between each active electrode pair was then calculated, a high pass filter with a 10 Hz cut-off frequency applied and EMG signals rectified [128, 352].

4.3.4 Functional MRI data analysis

PFM analysis

Functional MRI datasets were motion corrected with rigid body registration to the mean of the functional scans, and the realignment parameters were visually assessed for task-related movements, particularly at the times of the finger tapping events. This resulted in two datasets from one subject (left-handed) being discarded due to significant motion during the execution of the self-paced finger

tappings. Therefore, 10 datasets from the remaining five subjects were corrected for physiological noise using RETROICOR [144]. In-brain voxels were then masked for further analysis. Detrending was performed to remove trends of up to 4th-order Legendre polynomials, and sine and cosine signals with one cycle over the scan duration. Each voxel time series was then normalized to the mean value of its baseline volumes to compute the percentage signal change. These steps were performed using AFNI (NIMH/NIH) [82].

Each dataset was then analyzed using the PFM method described in the Theory section (§4.2), with algorithms implemented using in-house software written in Matlab (The Mathworks Inc., Natick, MA). Necessary for the ridge regression deconvolution, the raw autocorrelation coefficients $r(i)$ were estimated from the baseline time points in a voxelwise basis considering the three-dimensional nearest neighbours ($K = 27$), although smaller kernels were used at mask boundaries. The raw autocorrelation coefficients were used by the LD algorithm to compute the AR model parameters for each candidate order and the AIC, AICc, MDL and MDLc criteria were investigated for model order selection. The final autocorrelation coefficients used to define Σ were computed from the AR parameters (Appendix A.4) and normalized by $\hat{r}(0)$.

In order to ease the visualization of the noise serial correlations, we computed for each voxel [369]

$$S_p = \frac{B}{1 + 2 \sum_{i=1}^{B-1} r(i)}, \quad (4.25)$$

where $S_p = B$ indicates uncorrelated or white noise as $r(i) = 0$ for $i = 1, \dots, B - 1$. Lower values of S_p implies higher degree of autocorrelation. Two spatial S_p maps were computed: an S_p map using the raw autocorrelation coefficients but only considering the coefficients until the first lag with $r(i) < 0$ in order to avoid including those coefficient with very high variance, and an S_p map using the estimated autocorrelation coefficients after fitting the optimal AR model. Histograms with the number of voxels per AR model order, and spatial maps of the AR model coefficients, the optimal model order \hat{p} and S_p were obtained for each criteria [280]. Based on a qualitative analysis of the maps and the histograms (see §4.4), the estimates corresponding to the MDLc criterion were chosen to model the autocorrelation for the deconvolution. Furthermore, a maximum model order of 3 was established for the MDLc criterion since higher orders were mainly located on large veins and CSF regions, as shown in the T_1 - and T_2^* -weighted anatomical images.

Subsequently, the ridge regression (RR) deconvolution was performed using (4.14) where the convolution matrix H was defined from the two gamma-variate HRF with standard SPM (FIL/UCL) parameters [130] sampled at the corresponding TR. In order to increase the speed of the algorithm, the RR deconvolution was performed in blocks of 50 scans for TR 2 s and 150 scans for TR 0.4 s with the regularization parameter, λ , computed from the first block. To validate this block-by-block implementation, the analysis was repeated for several block lengths, and block length was found to have negligible effect on the results provided that the duration of the block was longer than the HRF. Following the RR deconvolution, t -statistics were computed using Eq. (4.20) with the five closest in-plane voxels ($L = 5$) and each t -map was FDR corrected and thresholded at FDR $\alpha \leq 0.05$. Finally, spatial clustering was applied to the thresholded maps with a minimum three-dimensional cluster size of 5 voxels in order to enhance the spatial specificity of the method, i.e. diminish the probability of isolated false positives. The resulting t -maps were then overlaid on the T_1 -weighted and T_2^* -weighted anatomical images which were resampled to EPI space using a 12-parameter affine spatial transformation.

Activation movies, Activation Events and Activation Time Series

A problem with paradigm-free fMRI analysis is that the whole data set must be explored to determine when interesting cortical events occur. Two methods were developed to address this. To begin with, an activation movie was created from the time series of thresholded t -maps to allow an exploratory and qualitative visualization of when coordinated activation occurred across the cortex. Secondly, an Activation Time Series (ATS) was created to compress this 4D thresholded dataset, shown in the activation movie, into a 1D time series that highlighted periods when coordinated activation occurred across the cortex (see Figure 4.9). Each plot of ATS contains two lines to separate positive and negative BOLD events: the positive going line (black) plots the number of voxels exceeding the t -threshold with positive signal amplitude at each time point, whereas the negative going line (red) similarly plots voxels with negative signal amplitude.

In addition to the finger tapping events, all activation events during apparent periods of rest involving more than 100 voxels on the ATS were investigated. For each activation event, an additional PFM Statistical Map was created by condensing the corresponding t -maps into one volume, where the value for each voxel was the maximum of the t -maps during that activation event. The number of t -maps

included in each activation event was determined according to the width of the corresponding peak in the ATS. Besides, for the activation events associated to finger tapping trials, a PFM Delay Map displayed the latency between the time of the maximum statistically significant activation included in the 'PFM Statistical Map' and the movement onset recorded by the EMG. In addition, for those voxels declared active in each trial, the time series of percent BOLD signal changes were measured from 5 s before to 25 s after movement onset, and filtered with a Hamming window, low pass filter with cut-off frequency of 0.3Hz, to investigate the variability of the hemodynamic response across trials.

To summarize and combine the results for all subjects, the number of times a functionally interesting ROI exceeded threshold at an expected event for each TR and each type of trial (cued and self-paced) was found, in order to calculate its frequency of activation and investigate inter-trial consistency. For the paradigms used here, functionally interesting ROIs were identified following relevant publications [83, 368] and included the supplementary motor area, bilateral lateral premotor cortex, bilateral primary motor cortex, bilateral primary somatosensory cortex, superior parietal cortex, inferior parietal cortex, and secondary and primary visual cortex.

Evaluation of PFM

Three methods were used to evaluate the proposed PFM method:

- i) The EMG time-series of the hand movements were compared to the onset of activated regions by means of a nonparametric test of the correlation between the ATS and the EMG signals.
- ii) The spatial extent of activated regions was compared to the statistical maps obtained with the traditional GLM analysis approach, using the timing information from the EMG signals as the stimulus onsets for the definition of the regressors.
- iii) The spatial extent of activated regions was compared to the results of the Probabilistic Independent Component Analysis algorithm [27] included in FSL.

i) Rank correlation between ATS and EMG signals

The Spearman's rank correlation coefficient [171] between the the EMG amplitude time series of the LE, RE and RF and the fMRI ATS (positive and negative

lines summed) were calculated using the Matlab Statistical Toolbox (The Mathworks Inc., Natick, MA), along with the corresponding p -values for the non-correlation hypothesis. The EMG signals were first decimated to have the same number of time points as the ATS and then thresholded at an amplitude z -score of 4, i.e. 4 times the standard deviation from the EMG time course mean, in order to only consider significant hand movements. The decimation was performed using the Matlab Signal Processing Toolbox (The Mathworks Inc., Natick, MA) which applies a Chebyshev Type I filter with normalized cut-off frequency $0.8/r$, where r is the decimation factor, and 0.05 dB of passband ripple. The correlation coefficients were computed to assess the accuracy of PFM in detecting significant movements recorded in the EMG. Note that in case of no movement in the decimated time series, the Spearman's rank correlation cannot be defined.

ii) GLM analysis

A GLM-based analysis was done with the regressors in the design matrix calculated from the convolution of the stimulus time series for the VCTs and SPTs (modelled as delta functions at the times of onset of tapping as detected by EMG) with the SPM-canonical HRF and its first temporal derivative [130]. The GLM-based analysis was done using the 3dREMLfit function in AFNI which assumes an ARMA(1,1) model for the noise serial correlations, and the corresponding statistics (F-test) were FDR corrected and thresholded at FDR corrected p -value ≤ 0.05 and minimum cluster size of 5 voxels. GLM-based statistical maps were also computed for any activation events detected during periods of rest, but in this case using the onset times estimated from the ATS to define a single-event regressor.

In order to examine the degree of concordance between the GLM and PFM maps, we calculated the number of overlapping voxels in both maps for each of the tapping events ($\text{PFM} \cap \text{GLM}$), and the percentage of overlapping voxels relative to the number of detected voxels with PFM ($\% \text{ PFM in GLM} = \text{PFM} \cap \text{GLM} / \text{PFM}$).

iii) Probabilistic ICA

The datasets were also analyzed with the Probabilistic ICA [27] algorithm available in FSL (MELODIC, www.fmrib.ox.ac.uk/fsl). This software first performs principal component analysis (PCA) in combination with a model selection criterion, based on a Laplace approximation to the posterior distribution of the model evidence (LAP) in order to reduce the dimensionality of the data. The spatial independent

components, or sources, are then computed with the FastICA algorithm [177], which maximizes the non-Gaussianity of the spatial sources. Finally, the spatial maps are transformed into z -score maps by dividing by the voxelwise estimated standard error of the residual noise [27] so that statistical inference for the voxels in each spatial component can be obtained. In order to identify relevant components, we computed the Pearson correlation coefficients between the time course of each spatial component and a reference time course, which was generated by the convolution of the decimated right-flexor EMG signal with the two gamma-variate HRF with SPM default parameters (EMG-regressor). Spatial components with correlation coefficient larger than 0.3 were labelled as relevant components. For illustration, the z -score maps were thresholded at $z > 3.0$ and minimum cluster size of 5 voxels.

4.4 Results

Figure 4.5a illustrates the pre-processed, (i.e. corrected for motion, trends and physiological fluctuations), normalized BOLD signal time course of a voxel located in the contralateral primary motor cortex for a representative subject at TR 2 s. Figure 4.5b shows the corresponding sequence of t -maps (p -value ≤ 0.05 , FDR-corrected) for the first visually-cued tapping (VCT) and the second self-paced tapping (SPT) produced using the PFM method. The timing of the maps are marked with red arrows in Figure 4.5a. The t -maps are shown for three different slices, for three consecutive time points (TR 2 s) labelled relative to the onset of movement as recorded by the EMG signals.

Statistical Modelling of the Noise Autocorrelation

Figures 4.6 and 4.7 depict the histograms of the optimal AR model order computed by the AIC, AICc, MDL and MDLc criteria considering all datasets at TR 2 s and TR 0.4 s, respectively. The maximum order shown in the histograms corresponds to the maximum order estimated with MDLc ($p = 3$ for TR 2 s, and $p = 10$ for TR 0.4 s) in order to facilitate the visualization, but the maximum model orders of AIC were $p = 4$ for TR 2 s and $p = 19$ for TR 0.4 s. The histograms clearly illustrate that slightly higher model orders are necessary to model the autocorrelation at TR 0.4 s due to the faster sampling rate. Besides, the histograms also demonstrate that optimal model orders are higher for the AIC-based criteria than for MDL-based criteria, and the finite sample criteria (AICc and MDLc) select lower model orders

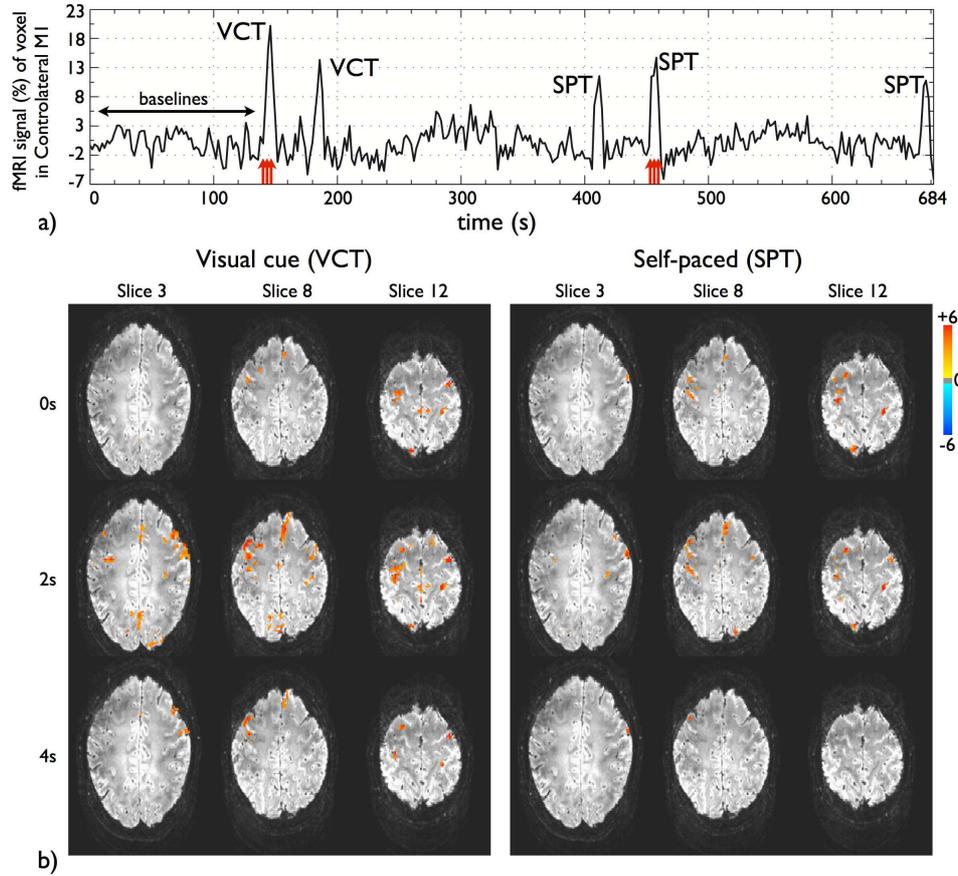


Figure 4.5: Data for subject A and dataset of TR 2 s (a) pre-processed and normalized fMRI time series of a voxel located in the contralateral primary motor cortex. The figure illustrates the initial baseline of B time points used to estimate the noise covariance matrix and to compute the statistics; (b) time course of t -maps of the first Visual Cued Tapping (VCT) and second Self Paced Tapping (SPT) indicated by the red arrows in a). The times shown on the left are relative to the movement onset as recorded by the EMG signals. Activity is detected (p -value < 0.05 , FDR-corrected) in SMA, bilateral PM, M1 and S1, inferior and superior parietal cortex and visual cortex.

than the asymptotic ones (AIC and MDL) owing to stricter penalization terms (see Table 4.1 and Figure 4.2).

To investigate this further, Figure 4.8 shows the spatial maps of: a) S_p computed with the raw autocorrelation coefficients; b) S_p computed with the raw autocorrelation coefficients; c) optimal model order \hat{p} ; and the AR model coefficients for orders d) $p = 1$, e) $p = 2$ and f) $p = 3$, for the dataset of Subject A at TR 0.4 s. This figure illustrates that considerable spatial variability and anatomical structure exists in S_p , higher (low degree of autocorrelation) in WM than in GM regions (Figure 4.8a). Consequently, WM areas with low correlation were modelled with an AR(0) model. It was also observed that model orders larger than 3 were mainly located

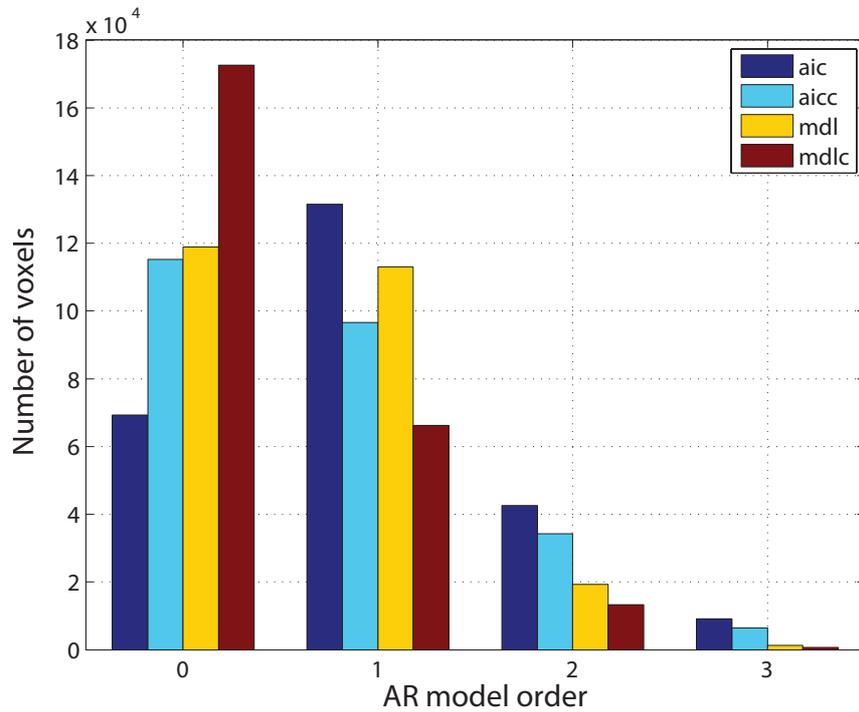


Figure 4.6: Histograms of the optimal model orders selected with the AIC, AICc, MDL and MDLc criteria for all dataset at TR 2 s (Number of baselines $B = 70$).

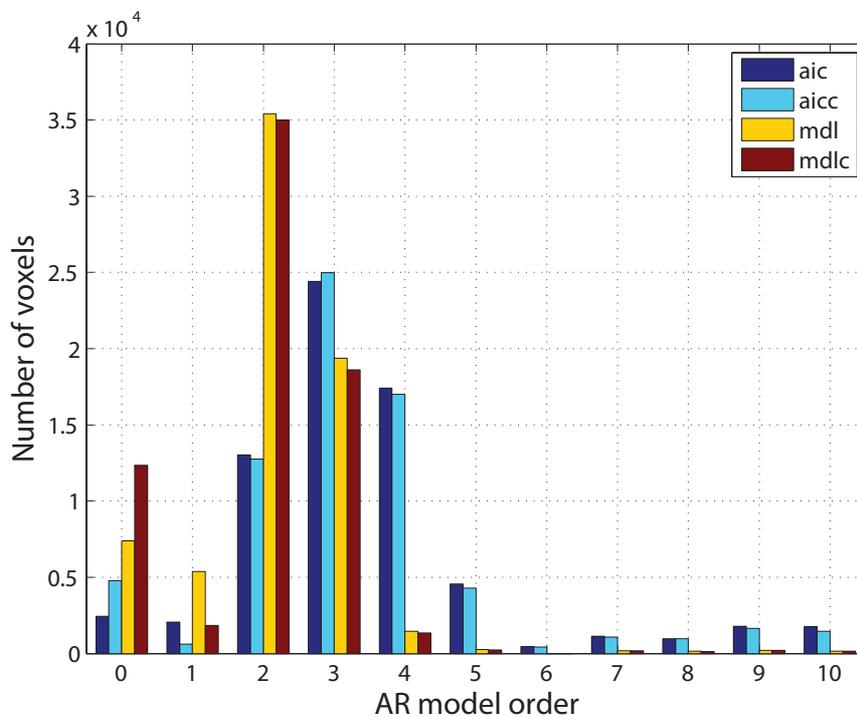


Figure 4.7: Histograms of the optimal model orders selected with the AIC, AICc, MDL and MDLc criteria for all dataset at TR 0.4 s (Number of baselines $B = 350$).

on not-masked CSF voxels and large veins as confirmed by the T_2^* -weighted image. This is shown in Figure 4.8c for the MDLc criterion, but was also observed for the other criteria. Regarding the AR parameter maps, the first two coefficients, a_1 and a_2 , are clearly more relevant for modelling the autocorrelation than a_3 , and even less significant for the coefficients of higher orders (maps not shown). In fact, for this dataset at TR 0.4 s, 90% of the voxels were modelled with orders $p \leq 3$. The first coefficient a_1 exhibited the largest amplitude (maximum value equal to -0.9)¹. Interestingly, a_1 provided very relevant anatomical information with larger values in GM regions, especially in areas of the posterior cingulate cortex, precuneus, lateral occipital cortex, central sulcus and superior frontal gyrus, as also shown in the S_p maps. Having said that, an AR(1) model is proven insufficient at TR 0.4 s by itself to model the autocorrelation, and we observed that the a_2 and a_3 coefficients also become meaningful for GM voxels. Exactly the same characteristics in the correlation were observed for the datasets at TR 2 s, where GM regions exhibited more correlation and so they were modeled with higher model orders. Considering these results, the MDLc criterion was considered for the deconvolution with a maximum candidate order of 3 which accounted for 96% of the voxels at TR 0.4 s and all voxels at TR 2 s, respectively.

Activation Time Series and EMG correlations

The single-trial detection sensitivity of the PFM method was assessed by comparing the fMRI ATS to the EMG signals. Figure 4.9 plots the RF and LE EMG signals and fMRI ATS for all datasets at each TR (the RF and RE EMG signals provided similar information). It can be seen that temporal clusters of activation detected as peaks in the ATS (without any knowledge of the EMG data in the fMRI analysis) are in synchrony with the finger tapping events or precede them by a maximum of 1 time point for the TR 2 s data and 4 time points for the TR 0.4 s data. In general, earlier activations were observed in the SMA and voxels located in the sulci of primary sensorimotor regions (pre-central and post-central sulci). The peaks shown in the ATS were used to compute the PFM Activation Maps where the number of t -maps depended on the length of the event and the TR of the dataset.

Although subjects were instructed to perform finger tapping only with the dominant hand, EMG activity was recorded in the left extensor in synchrony with the

¹Note that the definition of an AR model used in Eq. (4.17) implies negative values of the AR coefficients for positive autocorrelation

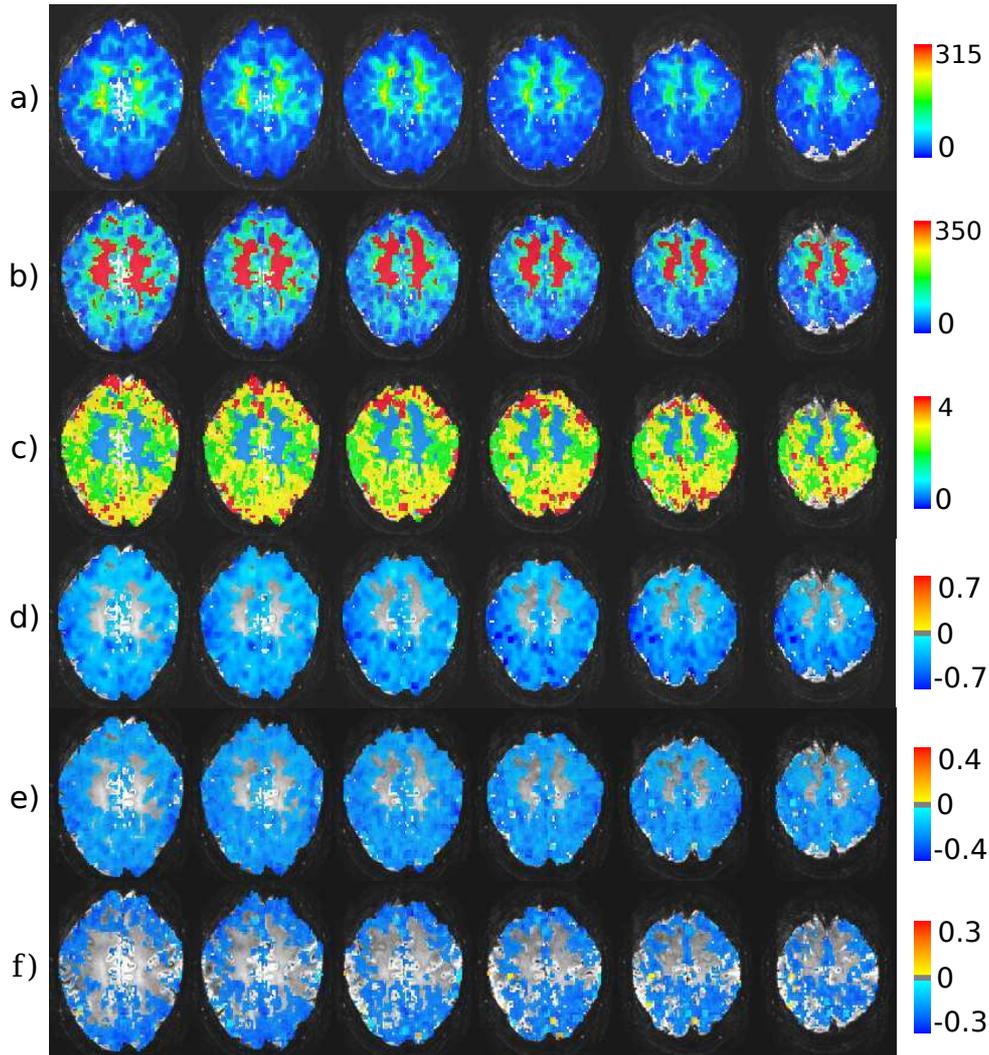


Figure 4.8: Spatial maps of a) S_p computed with the raw autocorrelation coefficients; b) S_p computed with the raw autocorrelation coefficients; c) optimal model order \hat{p} ; and the AR coefficients for orders d) $p = 1$, e) $p = 2$ and f) $p = 3$. These maps correspond to the MDLc criterion for dataset of Subject A at TR 0.4s.

right hand finger tapping in all but subject F (see Table 4.2 and Figure 4.9). In addition, two subjects did not exactly follow the instructions: Subject A performed an additional finger-tapping of 2 s at the end of the scanning period in the TR 2 s dataset, which was confirmed by post-scanning questioning and by EMG, whilst subject B performed four SPT trials rather than the instructed two trials for both TR 2 s and 0.4 s datasets (see Figure 4.9). The PFM method was able to detect 7 out of 10 (7/10) VCT events at TR 2s, 11/13 SPT at TR 2s. PFM did not detect the second VCT and the second and third SPTs of subject B, and the two VCTs of subject C. As a comparison, no voxels were detected with GLM ($Q < 0.05$, FDR-corrected) in the first VCT of subject C, whereas only 11 voxels exceeded this

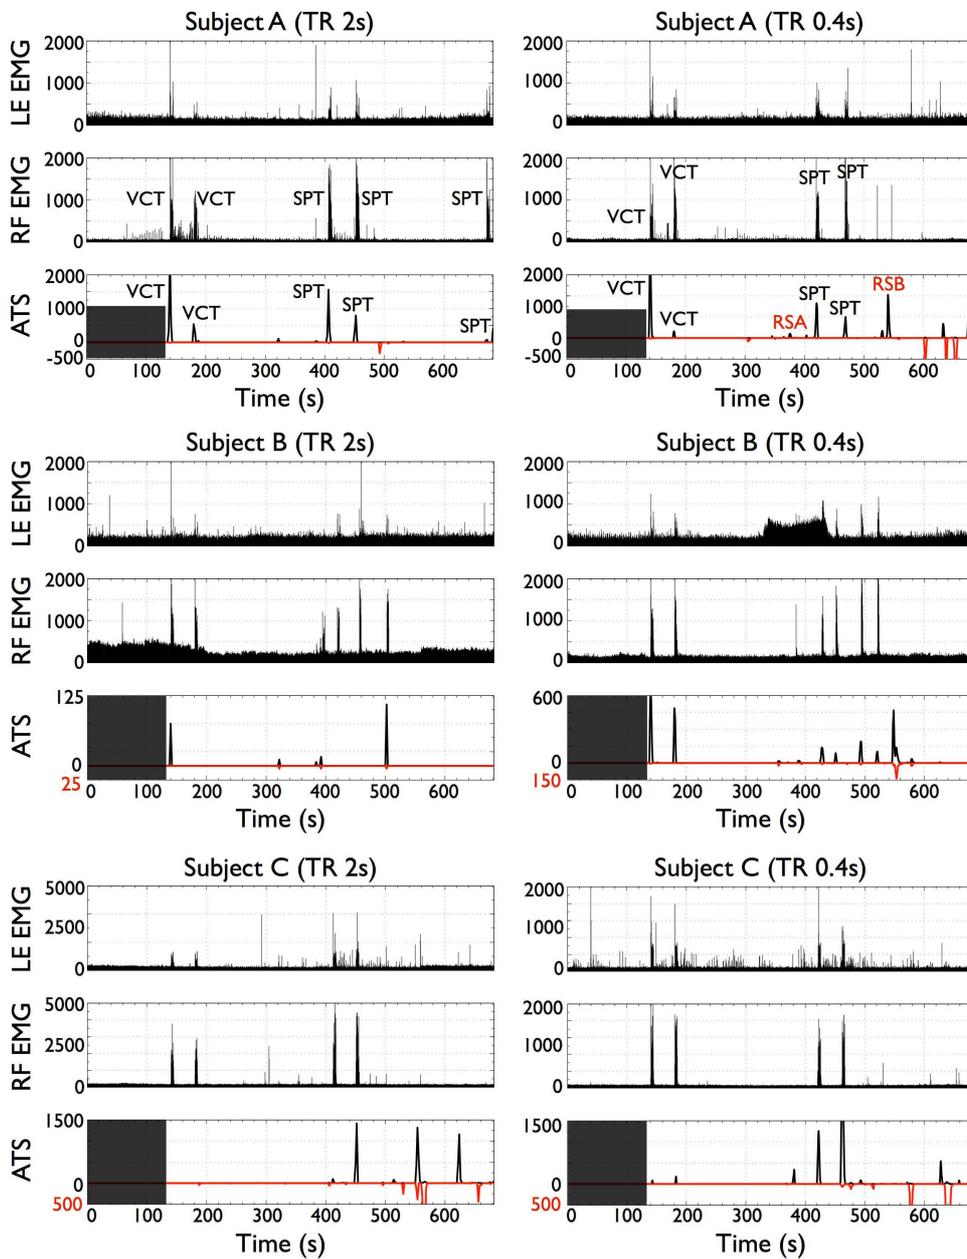


Figure 4.9: Activation Time Series (ATS) and EMG signals. Each plot shows the fMRI Activation Time Series ($Q < 0.05$, FDR-corrected) (bottom) and the EMG time series corresponding to the left extensor (LE) (top) and right flexor (middle) in order to capture upper limb movements. In order to facilitate the interpretation of the result, the ATS differentiates between positive activations (black, positive y-axis) and negative activations (red, negative y-axis). The dark box indicates the baseline period. Activation Time Series were used to detect periods of significant brain activation.

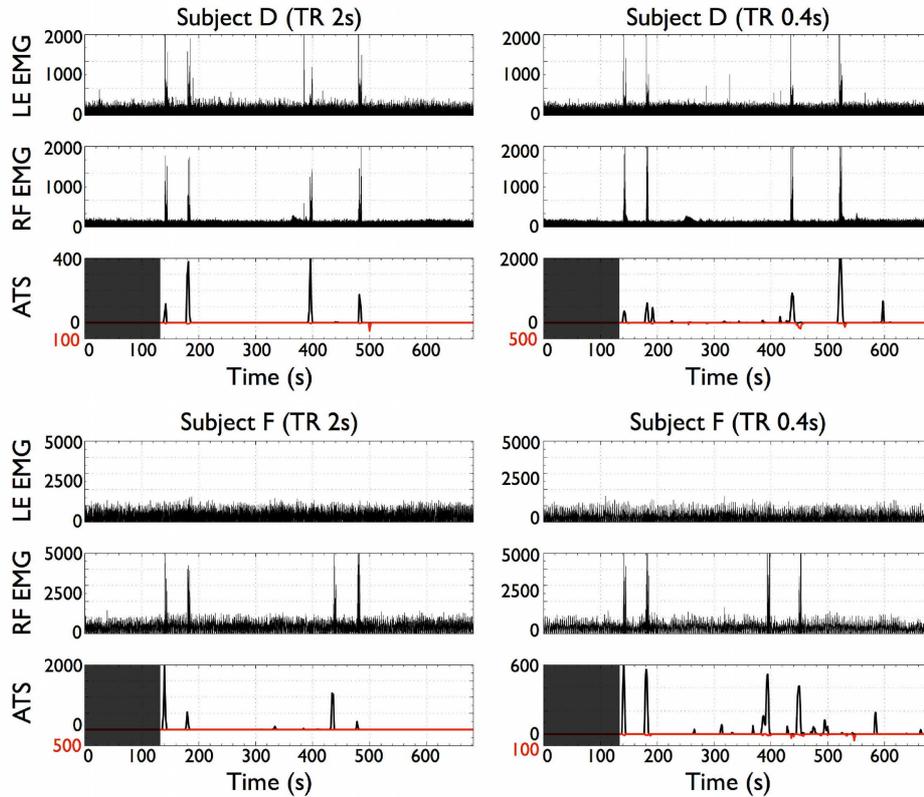


Figure 4.9: Activation Time Series (ATS) and EMG signals (Continued).

threshold in the second VCT. On the other hand, all finger tapping events were detected with PFM at TR 0.4s. Along with the VCT and SPT finger tapping events, all EMG measurements showed sporadic spikes of muscle activity in both hands which at times correlated with peaks on the fMRI ATS. Furthermore, the ATS also showed task-unrelated events, i.e. sporadic, spatio-temporal clusters of cortical activation occurring during the rest periods and which did not allineate with significant changes in the EMG signals.

Quantitatively, the concurrence between EMG activity and the activity detected with PFM was also evaluated by computing the Spearman's rank correlation coefficients between the fMRI ATS and the EMG signals. At both TRs, the fMRI ATS and EMG time courses were significantly correlated ($p < 0.001$) at 23 of the 30 possible combinations (EMG-dataset). The causes for no correlation between the fMRI signal and the EMG were: no detection of the event (RE and RF of subject C at TR 2 s); artifactual EMG acquisition (LE of subject B at TR 2 s); and reduced or no significant movement recorded by the EMG (LE of subject C at TR 0.4 s, LE of subject F at TR 2 s, LE and RE of subject F at TR 0.4 s).

	Left Extensor	Right Extensor	Right Flexor
Subject A	0.342*	0.483*	0.417*
Subject B	0.110	0.444*	0.262*
Subject C	0.242*	0.100	0.101
Subject D	0.620*	0.645*	0.444*
Subject F	-0.002	0.273*	0.446*

a) TR 2 s

	Left Extensor	Right Extensor	Right Flexor
Subject A	0.175*	0.326*	0.269*
Subject B	0.138*	0.251*	0.377*
Subject C	-0.011	0.103*	0.342*
Subject D	0.341*	0.427*	0.382*
Subject F	NaN	NaN	0.133*

b) TR 0.4 s

Table 4.2: Spearman’s rank correlation coefficients between the decimated and thresholded EMG signals and the Activation Time Series for datasets of TR 2s with 342 time points (a) and TR 0.4s with 1710 time points (b). The asterisks indicate that the correlation coefficient is statistically significant (non-parametric $p < 0.001$).

Comparison with GLM and trial-by-trial variability

Figure 4.10 illustrates the time course of t -maps ($Q < 0.05$, FDR-corrected) during the first VCT trial and the first SPT trial for one slice of subject A at TR 0.4 s, overlaid onto the corresponding T_2^* -weighted anatomical image. The duration of this sequence was determined from the corresponding peaks in the ATS (see subject A at TR 0.4 s in Figure 4.9). The PFM statistical map created by condensing the time course of t -maps with the voxelwise maximum is also shown and labelled PFM. The time shown in the t -maps is relative to the onset of movement execution as recorded by the EMG signals. For comparison, the statistical map (F-test, $Q < 0.05$, FDR-corrected) using the traditional GLM approach is also depicted in a box at the bottom right of each time sequence of t -maps in Figure 3. It can be observed that both PFM and GLM depict activations in overlapping regions. Averaging across all finger tapping events and all subjects (except the first VCT of Subject C at TR 2 s where no activated voxels were found by any method), we observed that 70% of the voxels identified as active in the PFM maps were also present in the GLM maps, with higher rate at TR 0.4 s (75%) than at TR 2 s (64%). The average value of % PFM in GLM at TR 2 s increased to 83% if the five events not detected with PFM at TR 2 s were discarded in this comparison.

Activations associated to the finger tapping tasks were found in areas of senso-

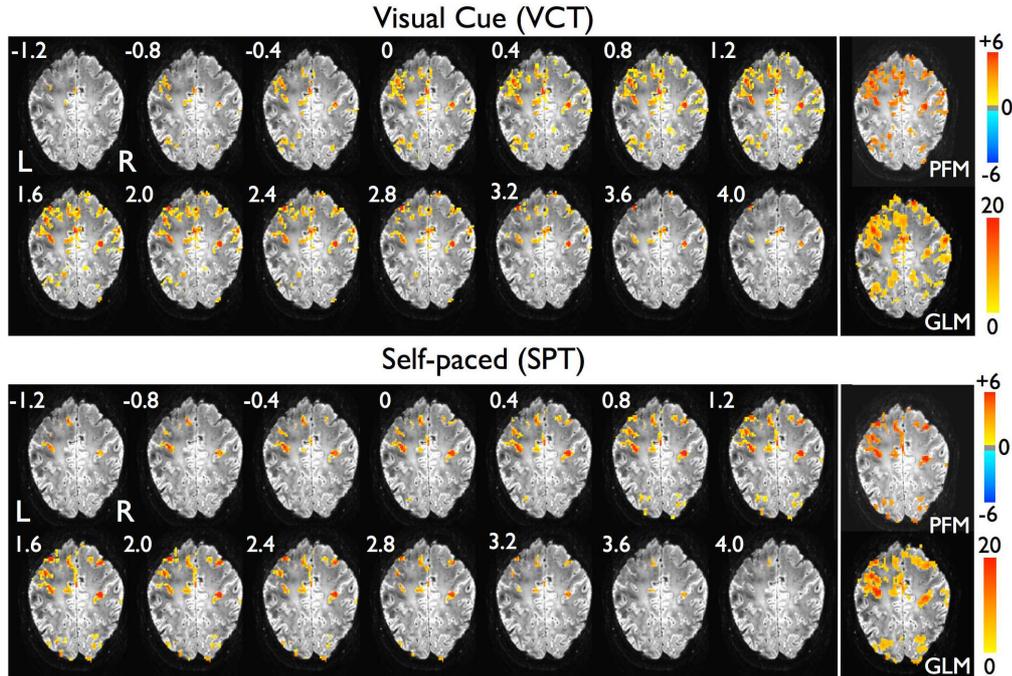


Figure 4.10: Time courses of t -maps ($Q < 0.05$, FDR-corrected) showing the activation for the first VCT (top) and first SPT (bottom) for subject A and TR 0.4s. In addition, the corresponding PFM Statistical map with the maximum t -statistic during the sequence and the GLM statistical parametric maps (F-test, $Q < 0.05$, FDR-corrected) after fitting a GLM model with the canonical HRF and the first derivative are shown at the right of each slice sequence (white box). The times are relative to the onset of the tapping as recorded by the EMG signals.

rimotor execution and processing of motor movements, as shown in the PFM maps in Figures 4.5b and 4.10. The relative timing of the different cortical responses showed large variability between trials. For the visually-cued tap shown in Figure 4.10, initial activity occurred in regions close to the contralateral central sulcus extending into the contralateral primary motor and primary somatosensory cortices, ipsilateral postcentral sulcus, ipsilateral supramarginal gyrus and bilaterally in the lateral occipital gyrus and intraparietal sulcus. Later, activation extended into a larger network of areas including the cingulate gyrus and posterior SMA (paracentral lobule), bilateral primary sensorimotor (M1 and S1) cortices, superior frontal gyrus and superior parietal lobule. Posteriorly, activation was also seen in the ipsilateral primary motor and bilateral premotor cortices, along with draining veins. Activation was found in similar regions for the self-paced tap shown in Figure 4.10, although ipsilateral premotor areas showed earlier activation in the self-paced tap than for the visually-cued tap. Interestingly, BOLD responses were seen in areas of

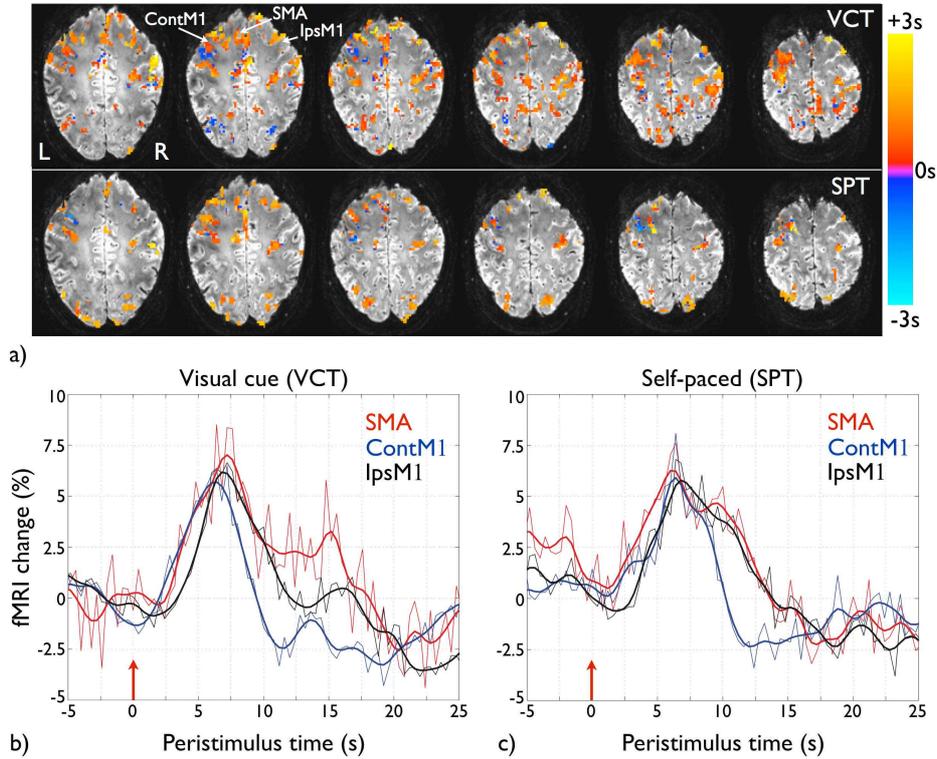


Figure 4.11: (a) Delay maps of voxels exceeding threshold ($Q < 0.05$, FDR-corrected) for each trial for the same slices shown in Figure 4.10. The delay is computed as the time when the maximum statistically significant activation occurred, defined relative to the movement onset recorded by the EMG signals; (b) Hemodynamic responses for the three highlighted voxels in SMA, contralateral M1 (contM1) and ipsilateral M1 (ipsM1) for the first VCT and (c) for the first SPT of the dataset of subject A at TR 0.4 s. The red arrows at time 0 s indicate the movement onsets as recorded by the EMG signals.

the occipital lobe not only in the visually cued task but also in the self paced trial.

Figure 4.11a shows the PFM delay maps for voxels exceeding the threshold ($Q < 0.05$, FDR-corrected), exhibiting large variability in the timing of the response across voxels both within and between functional areas, and across trials. Figures 4.11b and 4.11c illustrate this variability in the single trial hemodynamic response of the first VCT and the first SPT for three individual voxels located in the posterior part of the contralateral M1, contralateral SMA and anterior part of ipsilateral M1 (pointed in Figure 4.11). These voxels were selected to avoid any veins (detected by low signal in the underlying T_2^* -weighted image).

Both statistical and delay maps (see Figures 4.5b, 4.10 and 4.11) demonstrate variability in the spatial extent and delay of activation across individual trials. Variability was also observed between subjects. Table 4.3 lists the frequency of detection of events across subjects and trials, for both VCT and SPT tasks at both TRs.

	TR 2s		TR 0.4s	
	VTC	SPT	VTC	SPT
SMA	0.7	0.7	0.9	1
Contralateral PM	0.7	0.35	0.9	0.66
Contralateral M1	0.7	0.7	0.9	1
Contralateral S1	0.7	0.7	0.9	1
Contralateral Inferior Parietal	0.7	0.4	0.9	1
Contralateral Superior Parietal	0.56	0.53	0.72	0.75
Ipsilateral PM	0.56	0.53	0.72	0.75
Ipsilateral M1	0.7	0.6	0.9	1
Ipsilateral Inferior Parietal	0.56	0.53	0.9	1
Ipsilateral Superior Parietal	0.5	0.47	0.8	0.75
Primary Visual	0.45	0.53	0.75	0.83
Secondary Visual	0.56	0.4	0.8	0.5

Table 4.3: Frequency of activation in regions of interest (ROI). This value is defined as the number of events the ROI was found active divided by the number of times (tasks and subjects) for which this area was scanned.

Datasets showed consistent activations (frequency of activation > 0.70) across subjects and tapping tasks in the SMA, bilateral M1 and bilateral S1 regions. It was also observed that contralateral M1 and S1 areas showed the most repeatable patterns of activation across trials (i.e. same voxels consistently activated), followed by ipsilateral S1 and M1. It should be noted that coverage of the lateral PM in both hemispheres was limited by the tilt of the slices, especially for the datasets of TR 0.4 s which had fewer slices. In addition, the primary visual cortex was not covered in subject B for both TRs, and in subject C for TR 0.4 s.

Comparison with Probabilistic ICA

Figure 4.12 illustrates the six spatial probabilistic ICA maps selected as relevant (i.e. correlation coefficient with EMG-regressor > 0.3) and their associated time series for the same dataset as Figure 4.10. Probabilistic ICA is also able to map the cortical activation in areas of the primary sensorimotor cortex, although the temporal information provided directly by PFM is spread across several spatial ICA components and associated time courses. Moreover, some areas are simultaneously active in several components, e.g. voxels of the contralateral postcentral gyrus in components 2 and 3.

The total number of components identified from probabilistic ICA using the LAP criterion, the percentage of the variability of the dataset described by these components, and the number of relevant components, (i.e. those whose correlation

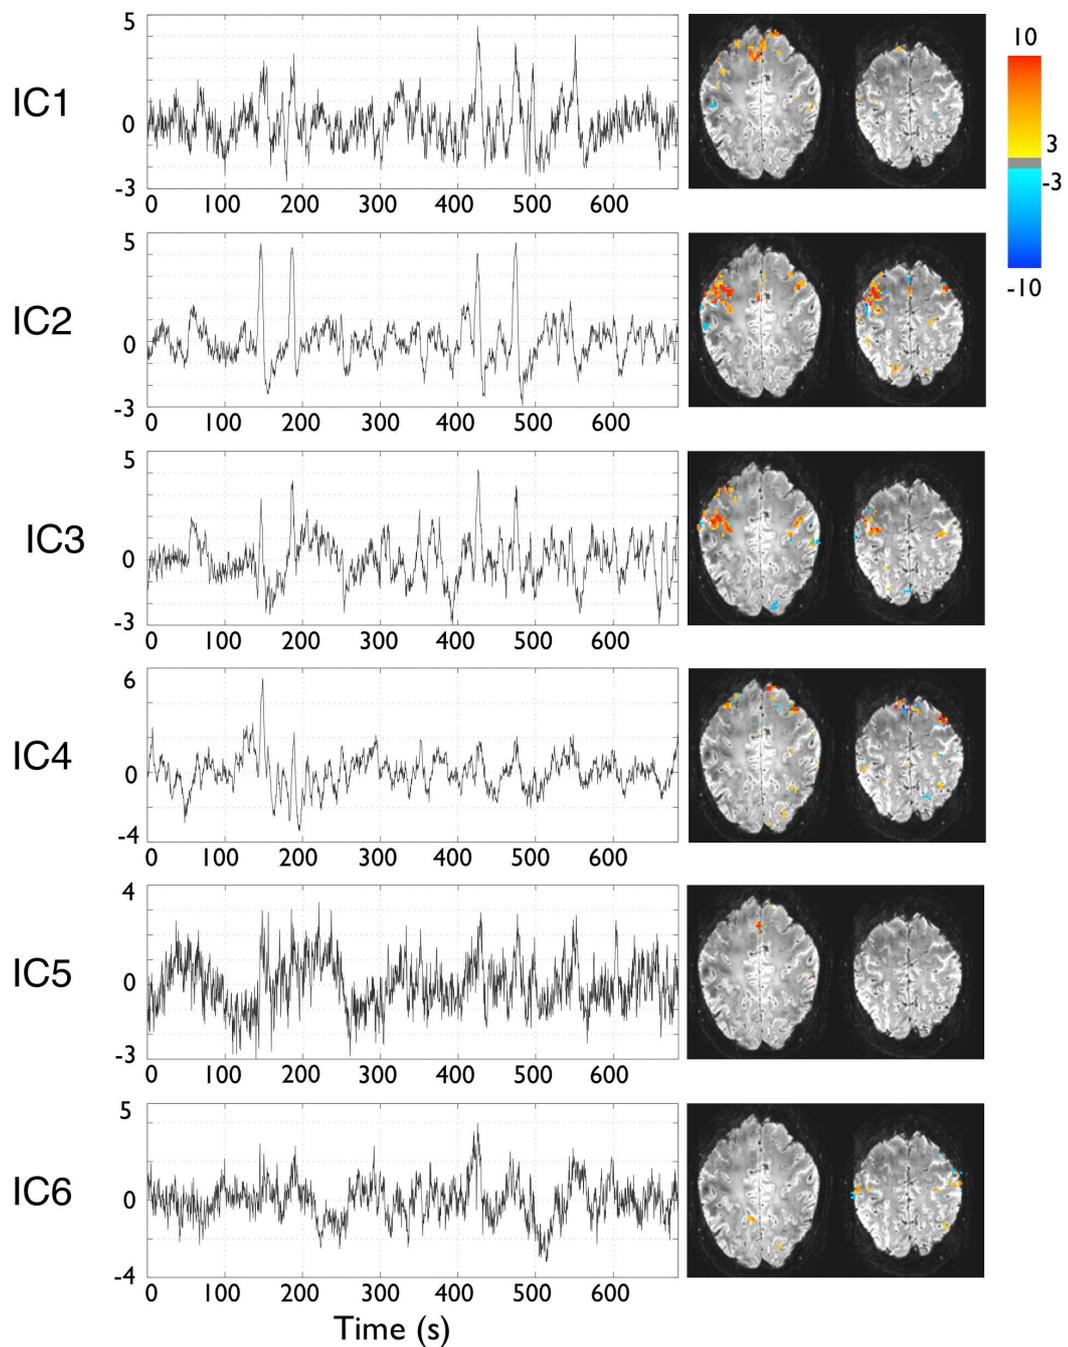


Figure 4.12: Spatial probabilistic ICA maps and associated time courses of the relevant components (i.e. those whose correlation coefficient with EMG regressor is > 0.3) for the dataset of subject A and TR 0.4 s (same as Figure 4.10 and 4.11). All maps represent z -scores (thresholded at $z > 3.0$) after normalization of the spatial components with the voxelwise estimated standard error of the residual noise.

	TR 2s		TR 0.4s	
	# ICs (% Var)	# Relevant ICs	# ICs (% Var)	# Relevant ICs
Subject A	134 (82.0)	11	431 (72.3)	6
Subject B	119 (79.9)	2	298 (63.7)	6
Subject C	124 (82.8)	12	556 (78.3)	33
Subject D	125 (82.6)	11	549 (77.0)	16
Subject F	113 (78.2)	2	561 (75.6)	4

Table 4.4: Number of Independent Components (ICs) identified with Probabilistic ICA using the criterion with maximizes the Laplace approximation of the posterior distribution of the model evidence (LAP), the percentage of the variability of the data retained with these components, and the number of relevant ICs (correlation coefficient with EMG regressor > 0.3).

coefficient with the EMG-regressor > 0.3) are shown in Table 4.4. The number of independent components revealed with Probabilistic ICA using the LAP criterion was very large for the datasets acquired with high temporal resolution (TR 0.4 s). There was considerable variability between subjects in the number of components identified as relevant. For the components shown in Figure 4.12, the ICA responses were more unilateral than the responses detected with either PFM or GLM, and the EMG indicated that there was actually movement of both hands (except for one subject). In fact, ICA also detected bilateral activation in the motor cortex as illustrated in components IC2 and IC3. However, the figure only plots those components that are highly correlated with the EMG-regressor (correlation coefficient > 0.3), suggesting that the correlation between the components including bilateral activations were not highly correlated with the EMG.

Transient task-unrelated activations

Figure 4.13 illustrates the PFM Statistical Maps ($Q < 0.05$, FDR-corrected) and the GLM analysis (F-test, $Q < 0.05$, FDR-corrected) for two task-unrelated activations (at points RSA and RSB in Figure 4.9) detected during the periods of rest for the dataset of subject A and TR 0.4 s. The GLM statistical maps were computed using the time of the peak in the ATS to define the onset time for these task-unrelated event. There was excellent agreement between the GLM statistical parametric maps and the PFM Statistical Maps. These task-unrelated activations did not show a specific recurring pattern or a consistent spatio-temporal pattern between subjects. For instance, the PFM map of event RSA showed cortical activity in the posterior cingulate cortex, precuneus, and lateral inferior parietal cortex,

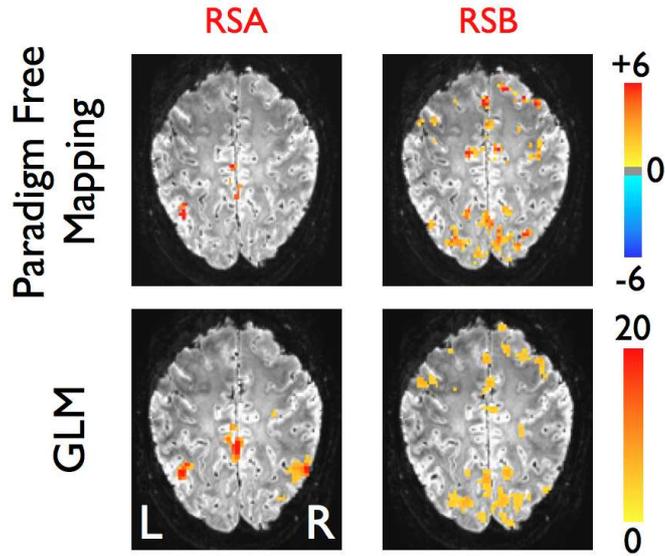


Figure 4.13: Statistical maps obtained with PFM ($Q < 0.05$, FDR corrected) (top) and GLM (F-test, $Q < 0.05$, FDR-corrected) (bottom) corresponding to the activations events shown at the ATS of subject A at TR 0.4 s in Figure 4.9.

whereas the map at point RSB showed activation in both primary sensorimotor areas (PM, M1 and S1, anterior cingulate cortex and SMA) and visual areas, and is related to task-independent hand movement (confirmed by the EMG).

4.5 Discussion

We have described a new Paradigm Free Mapping method for the analysis of fMRI data and shown that using this method it is possible to detect true single-trial BOLD responses to visuomotor activity at ultra-high field without any prior knowledge of the paradigm timing. This method differs from other totally or partially paradigm free techniques or other single trial analysis methods (see §4.1) in that PFM combines both concepts, enabling the observation of the dynamics of the BOLD responses of single trial responses in a simple and unsupervised manner with no knowledge of the stimulus timing.

4.5.1 Methodological issues

Signal model

The PFM method is based on the voxelwise deconvolution of the HRF under a linear model using the ridge regression algorithm and characterization of the noise

serial correlations as an AR stochastic process. This contrasts with the standard way that deconvolution is understood in fMRI analysis, which corresponds to the estimation of the HRF shape assuming prior knowledge of the paradigm [143, 149]. This is usually known as Finite Impulse Response (FIR) analysis. Here, we follow a more signal processing or linear systems perspective where the goal is to reverse the effect of the hemodynamic response over the neural-related signal [223]. Although rarely applied, this perspective has been proven adequate to study the effective and functional connectivity of neural responses with fMRI [142].

Ridge regression deconvolution

As shown in Appendix A.1, the L_2 -norm regularization used in the RR deconvolution is equivalent to a Bayesian deconvolution assuming Gaussian priors for the noise and the signal coefficients [142]. The use of Gaussian priors has been previously proposed in FIR approaches to compute smooth estimates of the shape of the HRF by setting a model for the correlation of the unknown coefficients, which is equivalent to the Tikhonov regularization [70, 78, 149, 241]. Here, by using the RR deconvolution, we assume uncorrelated coefficients as shown in Eq. (1). A Gaussian correlated prior for the coefficients could be used in PFM, but it is difficult to specify the correlation of the underlying signal in advance. Nevertheless, setting Laplacian priors which is equivalent to L_1 -norm regularization estimators [339, 340, 367] is more reasonable, and this is investigated in the following chapters.

Importantly, any PFM deconvolution model will rely on a particular shape of the HRF. Here, we used a two-gamma variate HRF with the standard SPM parameters, which includes a post stimulus undershoot [130]. Even though the ridge regression estimation has previously demonstrated high robustness against model mismatches within a limited uncertainty region [317], excessive differences in the shape of the HRF could introduce bias in the deconvolved signal, either in timing and amplitude, and in turn the t -statistics. It would be interesting to determine the effect of variation of the HRF model on the results and compare this information with existing models of the variability of the HRF [4, 103, 162].

Statistical modelling of the noise autocorrelation

The parameters of the AR model of the noise serial correlations were estimated with the Levinson-Durbin algorithm and using the time points during a baseline or pre-stimulus period. Therefore, contrary to the traditional procedure in which

the autocorrelation function is estimated from the residuals after an initial simpler fitting, our approach directly estimates the autocorrelation from the 'noisy' periods of the time series and the estimate is not affected by the deconvolution matrix \mathbf{H} [373]. A drawback of this approach is that the estimated autocorrelation model could be inaccurate if the noise autocorrelation changes throughout the acquisition. Non-stationary models for the noise could also be investigated using a PFM formulation [225].

The datasets were corrected during preprocessing for motion, physiological cardiac and respiratory-related fluctuations with RETROICOR [144], and sinusoidal and polynomial trends up to 4th order. This helped to reduce the most relevant low-frequency fluctuations and facilitated the use of low-order AR models contrary to high-order AR or $1/f$ noise models studied in [56, 135, 382].

The raw autocorrelation coefficients were estimated for each individual voxel considering a three-dimensional neighbourhood. Spatial averaging of the autocorrelation coefficients considerably reduced their variability [373]. Next, the raw autocorrelation function was fitted to an AR model of variable order as it has been proposed in [55, 280, 373]. Our results demonstrated considerable spatial variation and anatomical dependence of the autocorrelation, with higher correlation on gray matter than white matter areas as shown in the spatial maps of S_p and optimal model orders. These characteristics of the autocorrelation agree with the results in Woolrich et al. (2001) [369] and Worsley et al. (2002) [373], but seem to contradict the results by Penny et al. (2003) [280] which observed more correlation in white matter regions. This difference in results may be explained due to differences in the magnitude of the MR fields and the ratio of physiological to thermal noise. In Penny et al. (2003) the analysis was done on datasets acquired at 2T, whereas a 3T system was used in Woolrich et al. (2001) and a 7T system is used in our experiments. Since the ratio of physiological noise to thermal noise increases with the static magnetic field and it is more prominent in GM [200, 201, 343], it is expected that the higher AR orders necessary to model low-frequency physiological fluctuations prevail over zero orders of thermal noise. Although a plausible explanation (also hinted in [56]), the anatomical description of temporal correlations might also depend on MR acquisition parameters, such as TR, TE or voxel size [44, 343], along with physiological, hemodynamic and metabolic fluctuations in gray matter.

We investigated four model order selection criteria to select the optimal order of the AR process: AIC, AICc, MDL and MDLc. The histograms and spatial maps

confirmed that the AIC-based criteria tended to select larger AR orders and overfit the noise. It must be noted that the model selection criteria based on the variance of the reflection coefficients or partial correlation coefficients [46] have been proven to give similar performance to the AIC criterion [49], and this could be the reason why AR models performed badly in [369]. In addition, using the LD algorithm enabled us to compute the AR model parameters for multiple candidate orders in a faster and more efficient way than using Variational Bayes approaches [280]. The evaluation was qualitatively performed based on the uniformity of the spatial maps of the autocorrelation and model orders across tissues, and more compact histograms (e.g. unimodal distribution of the AR model order). This allowed us to observe that a maximum AR order of 3 seems to be sufficient to model the autocorrelations in gray matter voxels, with higher orders located in CSF voxels or large veins, in agreement with the histograms of model orders in [56, 280, 369].

Temporal t -statistic and multiple hypothesis correction

The statistical significance of the activations was evaluated after FDR thresholding of the temporal t -statistic. Similar to Lindquist et al. (2007) [215], the t -statistics were computed as the deviance of the estimated signal time series for each voxel from its estimated baseline signal. However, the PFM t -statistics are based on filtering the fMRI time series with a filter based on the deconvolution of the HRF, whereas in Lindquist et al. (2007) the t -statistics are computed after applying an exponentially weighted moving averaged (EWMA) filter to the voxel time series. Consequently, there is a noticeable difference between the interpretation of the EWMA statistics and PFM statistics since the latter is directly related to the signal underlying the BOLD response.

In any event, setting baseline periods involves making the assumption that no hemodynamic event of neural origin occurred in those periods and our own results have illustrated the presence of significant and coherent activations during periods of rest different from the initial baseline. In a few datasets this might have lead to loss of sensitivity to some responses in those areas with significant variability during the baseline period [284]. All in all, the baseline period needs to extend a sufficient number of scans so that accurate estimates of the mean baseline amplitude and the variance are obtained. We used large values of B , but in practice we could use fewer baseline time points since the Student's t -distributions does not change considerably providing B is large enough. The baseline period also needs to be of sufficient length

in order to potentially include several cycles of baseline fluctuations.

The Benjamini-Hochberg (BH) procedure was used to control FDR for each of the time points separately [32]. No temporal correction was performed because the spatial and temporal correlation of the statistics led to a significant reduction in the spatial (but not temporal) sensitivity of PFM to the main events. Future work should address the investigation of spatio-temporal cluster inference methods in a PFM formulation, e.g. based in adaptive FDR procedures [314] or current cluster thresholding methods [77, 325]. These procedures will account for the spatial and temporal smoothness of the PFM statistics and increase the sensitivity of the technique in a four-dimensional correction. Furthermore, the sensitivity and specificity of any statistical correction method based on FDR could be enhanced by limiting hypothesis testing to voxels located within cortical layers [122, 344].

Activation Movies and Activation Time Series

The output of the PFM analysis method is a time course of t -maps which displays the dynamics of cortical activations in space and time. To facilitate an exploratory analysis of the four-dimensional (space and time) results, creating a movie of the activation proved very useful. Brain activation movies have previously been employed to display the cortical response to stimulus processing [363], but based on a FIR model approach where the onsets of the stimuli were known a priori. In addition, an Activation Time Series (ATS) was generated for each dataset which plots the number of voxels revealed as active with PFM. This helped us to reduce the dimensionality of the data and make the results easier to interpret. Alternatively, we could plot the sum of the t -statistics above threshold. Similar procedures have been proposed in Temporal Clustering Analysis [220, 257]. Exploring the results in space, instead of in time, could be done by computing voxelwise statistics (e.g. F -statistics or R^2 -coefficients) summarizing the relevance of the detected activations in order to create statistical maps according to a given threshold [297, 35], and then examining the corresponding time courses to identify the timing of the events.

4.5.2 Experimental results

Correlation with EMG recordings

The performance of the PFM method was evaluated using a paradigm where subjects performed dominant-hand finger-tapping, either visually cued or self paced.

Nonparametric p -values based on the Spearman's rank correlation coefficient demonstrated significant correlation between the EMG signals and fMRI ATS, and both were able to detect visually cued and self paced motor movements. We observed a higher probability in the detection of the events at TR 0.4 s (22 out of 22 events) than at TR 2 s (18 out of 23 events) despite the lower CNR available at TR 0.4 s. This might be due to improved characterization of the HRF shape at high temporal resolution (TR 0.4 s). Nevertheless, further investigation is required to evaluate the effect of TR and CNR on the deconvolution. In some cases the peak on the ATS preceded the EMG peak by up to one time point for data with TR 2 s or four time points for the TR 0.4 s. Earlier onsets may have arisen due to cortical activity associated with preparation and planning, or may indicate mismatches between the canonical HRF and the actual HRF. Once the activations have been detected with PFM, a posterior fitting of the HRF parameters and their correlation with EMG features could be helpful to understand the origins of any difference and investigate within voxel differences in the response between different conditions, such as self-paced, cued and transient task-unrelated activations [30, 52, 214]. Importantly, the coupling between cerebrovascular responses (CBF, CBV and CMRO₂) and neural activity must be considered in order to make inferences about whether differences in timing correspond to vascular or neuronal changes [59].

Anatomical and functional correlates

The PFM technique did consistently detect significant activation in brain regions known to be involved in the initiation and processing of a motor task [368], including SMA [83, 84], bilateral PM [158, 165, 278], bilateral M1 [83, 298], bilateral S1 [290, 296], bilateral inferior parietal (supramarginal gyrus) [165], superior posterior parietal areas [146, 313]. In addition, activation was also observed in the intraparietal sulcus and superior frontal areas in some subjects and these areas may be involved in the processing of visuo-motor tasks and spatial attention [80]. Primary and secondary visual cortex activity was observed for VCT and interestingly, activation in the primary visual cortex was also observed during the SPT trials (see Figure 4.10 and Table 4.3), possibly due to the mental visualization of the visual cue associated with previous finger tapping events. The inter-trial consistency of activations across regions was measured as an indicator of the functional relevance of a brain region for processing or performing a task [364]. Table 4.3 indicates that primary sensorimotor areas (SMA, M1 and S1) were the most consistently active areas

within and across subjects, whilst sensory association areas not primarily implicated in the task demonstrated larger variability [103, 298].

Ipsilateral activation in the PM and M1 cortices was detected for all subjects (see Table 4.3). The EMG signals indicated that all subjects, except subject F, had significant movement of the non-dominant hand during the uni-manual finger tapping task. However, even in the case of this subject where no activity in the non-dominant hand was recorded by the EMG, activity in ipsilateral sensorimotor areas was detected, supporting indications that the ipsilateral hemisphere may also play an important role in the planning, execution and control of unimanual motor tasks [187, 289].

Comparison with GLM and Probabilistic ICA

For validation we compared the activations revealed with PFM with GLM and Probabilistic ICA [27]. As illustrated in Figure 4.10, the PFM activation maps showed large correspondence with the GLM maps, especially at the main clusters of activation. In general, the PFM areas were smaller than the GLM areas but lay with a high degree of overlap (%PFM area in GLM area: 75.5% for TR 0.4 s and 64.2% for TR 2 s) suggesting that the thresholds applied in PFM are stricter than those used in GLM, which seem to indicate a higher specificity of PFM compared to GLM. Therefore, PFM could serve as an exploratory sequential technique to detect single-trial events, which later can be assessed with standard GLM techniques to increase sensitivity based on the times shown in the ATS. Importantly, differences between the PFM and GLM activation maps might also result from differences in the analysis procedure, e.g. modelling of serial correlations (PFM is based on an AR(p) models, whereas GLM analysis done with AFNI is based on an ARMA(1,1) model) or the way spatial information is incorporated (spatial smoothing of the data versus spatial averaging of the statistics) [334].

On the other hand, both PFM and ICA successfully detected time points and areas of activation associated to the finger tapping responses without prior knowledge of the paradigm, but there are numerous differences between both approaches. To begin with, in order to detect single trial responses with ICA, prior temporal segmentation of the datasets or posterior segmentation of the components into the single-trial frames is necessary. Methodologically, PFM is based on the deconvolution of an assumed model of the HRF, whereas ICA is completely model-free. The component decomposition obtained with some ICA algorithms, such as Fast-ICA or

Infomax, is sensitive to their initialization [168], whereas this is not a problem for PFM. In addition, ICA requires the number of components to be extracted from the data to be determined, either manually or with a model selection criteria (LAP, MDL or AIC). One of the most compelling aspects of PFM is the observation of the spatio-temporally coordinated activation in the activation movie. In comparison, each ICA component describes the behaviour of a number of voxels acting in synchrony, and so voxels with slightly different time courses may end up in different components which need to be appropriately combined to recover the temporal information that is immediately available from PFM. For datasets with numerous time points, such as our datasets acquired at TR 0.4 s, ICA requires a large number of components to describe the variability in the data (see Table 4.4) and functionally significant information is spread across several components (e.g. see areas of the postcentral gyrus which are included in the components IC2 and IC3 of Figure 4.12). Here, the relevant independent components were selected based on the correlation between the ATS and the EMG signal. This allowed us to identify spatial components including the SMA and bilateral primary somatosensory areas, but there were some areas, such as superior parietal lobule or intraparietal sulcus, which were not detected in the components selected by correlation with the EMG, but which were significantly active with both PFM and GLM (compare Figure 4.10 and 4.12). This does not imply that there are no components to account for these responses, but finding them may be very difficult in a real paradigm free mapping scenario. In fact, the PFM activation time course might prove a useful signature with which to identify relevant components. It is worth noting that the current PFM method presented in this chapter only detects discrete, temporally sparse events, and therefore would not be able to characterize components low-frequency oscillations, e.g. in resting state data, which can be detected with other techniques, such as ICA or seed-voxel correlation [88, 124]. Baseline fluctuations of the fMRI signal can be easily modelled within a PFM formulation and this will be investigated in Chapter 6.

BOLD response variability

Once a response had been detected with the PFM method, the high CNR available at ultrahigh MR fields allowed the amplitude and timing of the HRF to be observed on a trial-by-trial and voxel-by-voxel basis with no temporal averaging (see Figure 4.11). Results have demonstrated that amplitudes and delays vary considerably not only between voxels of different sensory areas, but also within the

same area confirming previous observations [103, 286]. In addition to differences in perception or execution across trials, this variability could be due to coherent spontaneous activity [125] or underlying baseline fluctuations of non-neuronal origin which were not removed with temporal detrending prior to deconvolution [336].

Transient task-unrelated activations

The PFM method enabled the unsupervised detection of significant transient, task-unrelated spatio-temporal patterns of brain activity across the cortex during periods of apparent rest (Figures 4.9 and 4.13). These activation events could be related to unconstrained behaviour in the scanner, such as small body movements with no specific recurring pattern, or may relate to other mental tasks such as changes in attention [80] or somatic perception [119]. Using the ATS, the temporal locations of these events were determined, and these events could subsequently be assessed using a traditional GLM analysis, as shown in Figure 4.13. In addition, PFM may also help to identify spurious extreme fluctuations which would affect any statistical analysis unless noticed and the information about the task-unrelated events extracted using PFM will complement the information obtained with other analysis techniques of resting-state BOLD fMRI. Future work will focus on assessing the functional significance of these activations and studying the factors giving rise to these task-unrelated or spontaneous events and their effects on cortical networks [284].

4.6 Conclusion

This chapter introduced the principles of paradigm free mapping, an fMRI data analysis technique which enables the detection and characterization of single trial BOLD responses without prior timing information. The method is based on the deconvolution of the fMRI time series under a linear model and assuming an HRF shape. In this chapter, we investigated an L_2 -norm regularized estimator, ridge regression. Statistical inference was obtained by defining t -statistics in the deconvolved data against a baseline period. The practicability of PFM was evaluated with a visuomotor paradigm with single trial visually cued and self paced finger tapping events and our results have demonstrated that PFM allowed for the first time the detection of true single trial events without specifying their onset timing.

The following chapters will investigate the use of L_1 -norm regularized estimators in PFM methods.

Chapter 5

Sparse Paradigm Free Mapping

This chapter introduces a Paradigm Free Mapping approach based on sparse estimation techniques: Sparse Paradigm Free Mapping (SPFM). In this study, two L_1 -norm regularized estimators are investigated, the LASSO and the Dantzig Selector. An homotopy-based algorithm is employed in order to compute all feasible estimates, and model selection criteria and shrinkage thresholds are investigated to select the regularization parameter (§5.2). The proposed method overcomes the need to specify a baseline period which was required in the previous PFM approach and ease the deconvolution of the BOLD events with the use of sparse estimators. The performance of the method is first evaluated with simulated fMRI data. Next, the operation of the technique is tested in the single-trial visuomotor paradigm described in 4. In addition, preliminary results investigating the potential application of SPFM to detect interictal epileptic discharges that occur between epileptic seizures without using information about the timing of the seizures and a comparison with EEG-based GLM analysis are also presented (§5.3). Finally, methodological and experimental discussions are given in §5.4.

5.1 Theory

5.1.1 Signal Model

The signal model considered in SPFM is the same as the one presented in Chapter 4. Assuming a linear model for the fMRI signal, the voxel time series can be modelled as

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \boldsymbol{\epsilon}, \quad (5.1)$$

where \mathbf{y} , \mathbf{s} , and $\boldsymbol{\epsilon}$ are vectors of length N representing the fMRI voxel time series, the neural-related signal driving the BOLD component of the fMRI signal and the noise, respectively. The matrix \mathbf{H} is a convolution (Toeplitz) matrix of size $N \times N$ defined from a given HRF. Zero-mean uncorrelated Gaussian noise with variance σ^2 is assumed in this chapter, i.e. $\boldsymbol{\epsilon} \sim N(0, \sigma^2 \mathbf{I})$. The methods presented in this chapter could be extended to correlated noise by means of data prewhitening providing the noise serial correlations were accurately estimated [373].

5.1.2 Sparse Estimation: The LASSO and the Dantzig Selector

In event-related single trial fMRI experiments BOLD responses are assumed to be the result of sparse, transient changes in neuronal activity within a cortical region [223]. Therefore, very few observations of the neuronal-related signal \mathbf{s} will exhibit significant amplitude at the time of the BOLD events, whereas most of the coefficients can be assumed to have negligible, zero amplitude. In that regime the signal \mathbf{s} is described as sparse. The simplest measure of the sparsity of a vector \mathbf{s} is defined as its number of nonzero coefficients, i.e.

$$\|\mathbf{s}\|_0 = \text{card} \{s_i | s_i \neq 0\}, \quad (5.2)$$

where a vector \mathbf{s} of length N is considered sparse if $\|\mathbf{s}\|_0 \ll N$ [50].

Contrary to the ridge regression (RR) estimator considered in Chapter 4 which regularizes the least squares (LS) solution with an L_2 -norm term, sparse estimators attempt to regularize the problem by maximizing the sparsity of the estimates, i.e. minimizing the number of nonzero observations (L_0 -norm) of the signal \mathbf{s} :

$$\hat{\mathbf{s}}_{L_0} = \arg \min_{\mathbf{s}} \|\mathbf{s}\|_0 \quad \text{subject to} \quad \|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2 \leq \delta_0. \quad (5.3)$$

Due to the discrete and discontinuous nature of the L_0 -norm [50], this optimization problem is NP-hard such that there is no polynomial time algorithm that can closely approximate the optimum value of the objective function [7] and finding the optimal solution involves a combinatorial search of all possible sparse subsets of the columns of \mathbf{H} . Thereby, approximate solutions are usually obtained based on the L_p -norm of \mathbf{s} , i.e. $\|\mathbf{s}\|_p$ with $p \rightarrow 0$ [74, 367].

Sparse estimators based on the L_1 -norm are usually a good approximation which allow the use of convex optimization algorithms, thereby avoiding the presence of local minima [166]. Based on (5.3), an estimate of \mathbf{s} can be found by solving the

following optimization problem

$$\hat{\mathbf{s}} = \arg \min_{\mathbf{s}} \quad \|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2 \quad \text{subject to} \quad \|\mathbf{s}\|_1 \leq \gamma, \quad (5.4)$$

or the equivalent problem

$$\hat{\mathbf{s}} = \arg \min_{\mathbf{s}} \quad \|\mathbf{s}\|_1 \quad \text{subject to} \quad \|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2 \leq \delta, \quad (5.5)$$

where γ or δ are nonnegative real regularization parameters. The optimization problem (5.4) is known as the Least Absolute Shrinkage and Selection Operator (LASSO) estimator [339], whereas (5.5) is known as the Basis Pursuit Denoising (BPDN) algorithm with inequality constraints [76]. It can be shown that both algorithms converge to the same solution with a one-to-one correspondence between γ and δ [120].

If the constraints are incorporated into the optimization function via Lagrange multipliers, the minimization of the Lagrangian of these estimators can be defined as

$$\hat{\mathbf{s}} = \arg \min_{\mathbf{s}} \quad \|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2 + \lambda \|\mathbf{s}\|_1, \quad (5.6)$$

where λ is also a nonnegative real parameter. The solutions to problems (5.4), (5.5) and (5.6) do not have closed form expressions, but they must be computed iteratively by solving a quadratic program (QP) for the LASSO formulation (5.4), a quadratically constrained linear program (QCLP) for the BPDN formulation (5.5), or a bound constrained quadratic program (BCQP) for the Lagrangian formulation (5.6) [120].

As an alternative to the LASSO or Basis Pursuits estimators, Candes and Tao (2007) proposed the Dantzig Selector (DS) which estimates \mathbf{s} with the following optimization problem [68]:

$$\hat{\mathbf{s}}_{DS} = \arg \min_{\mathbf{s}} \quad \|\mathbf{s}\|_1 \quad \text{subject to} \quad \|\mathbf{H}^T(\mathbf{y} - \mathbf{H}\mathbf{s})\|_\infty \leq \delta_{DS}, \quad (5.7)$$

where the L_∞ -norm of a vector \mathbf{s} is defined as $\|\mathbf{s}\|_\infty = \max_i \{|s_i|\}$. The DS problem can be recast as a linear program (LP) and therefore be efficiently solved for each value of δ_{DS} [68]. Like the LASSO, the DS optimization problem can be rewritten as

$$\hat{\mathbf{s}}_{DS} = \arg \min_{\mathbf{s}} \quad \|\mathbf{H}^T(\mathbf{y} - \mathbf{H}\mathbf{s})\|_\infty \quad \text{subject to} \quad \|\mathbf{s}\|_1 \leq \gamma_{DS}, \quad (5.8)$$

and the corresponding minimization of the Lagrangian function

$$\hat{\mathbf{s}}_{DS} = \arg \min_{\mathbf{s}} \quad \|\mathbf{H}^T (\mathbf{y} - \mathbf{H}\mathbf{s})\|_{\infty} + \lambda_{DS} \|\mathbf{s}\|_1. \quad (5.9)$$

Hereinafter, the presentation of the methods will be done based on the optimization problems minimizing the Lagrangian (5.6) and (5.9), and therefore the constant λ is used to refer to the regularization parameter regardless of the optimization problem. The computed LASSO and DS estimates depend on the regularization parameter, which controls the tradeoff between the sparsity of the estimate and the amplitude of the residuals, either their maximum absolute value in the case of the DS or their sum of squares in the case of the LASSO. A strong link between both problems exists since $\mathbf{H}^T (\mathbf{y} - \mathbf{H}\mathbf{s})$ is proportional to the derivative of $\|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2$ with respect to the coefficients of \mathbf{s} . Consequently, it has been observed that in a sparse scenario where most of the entries in \mathbf{s} are zero, the LASSO estimator and the DS exhibit very similar behaviour in terms of prediction error [36], becoming identical when the covariates of the model are orthonormal [180]. Note that the regressors of the design matrix \mathbf{H} are not orthonormal.

The L_1 -norm of the LASSO and the DS forces some of the coefficients of \mathbf{s} to be exactly zero as λ increases, inheriting excellent interpretability features of subset selection where only the most relevant regressors are kept with nonzero weights at the solution. From (5.4) and (5.7), one can notice that the LASSO and DS solutions converge to the least squares (LS) solution when the penalty parameter λ is equal or larger than the L_2 -norm of the LS estimate. In addition, the null solution, $\mathbf{s} = \mathbf{0}$, becomes valid when $\lambda \geq \|\mathbf{H}^T \mathbf{y}\|_{\infty}$ and thus this value is usually employed for the initialization of λ in homotopy algorithms which efficiently compute the complete path of feasible solutions, called the regularization path [108, 272].

5.1.3 Solving l_1 -minimization problems

Numerous optimization algorithms and codes have been proposed to efficiently solve the LASSO and DS problems in their various optimization problems or minimization of the Lagrangian function (see [345] for an overview). In this study, we rely on homotopy continuation procedures which enable the computation of the entire regularization path for all nonnegative values of λ (Figure 5.1) [15, 108, 180, 272]. Homotopy algorithms rest on the idea that the coefficient path is piecewise constant as a function of λ [272]. Initializing with $\lambda = \|\mathbf{H}^T \mathbf{y}\|_{\infty}$ and the null solution $\mathbf{s} = \mathbf{0}$,

the algorithm only changes from one suitable solution to the next one at critical values of λ where the active set of nonzero coefficients is updated by either including or removing one coefficient. Since homotopy algorithms work at most with the coefficients in the active set, the cost of each step is equivalent to solving a least squares problem of equivalent size [15, 108]. These properties make homotopy procedures very computationally efficient with highly sparse signals. In addition, explicit computation of the residuals can be obtained at each iteration so that the algorithm can be stopped when the sum of squares of residuals is below some theoretical threshold.

In particular, the homotopy algorithm developed in [15] is used here to compute the LASSO and DS regularization path¹. The algorithm is based on the strong duality between the primal and dual problems of (5.7) and (5.4) [47]. Since each step of the Primal-Dual Pursuit algorithm only requires one-rank matrix operations to find the update directions for the primal and dual updates, this algorithm is more efficient than other homotopy algorithms, such as the DASSO [180] or LARS algorithms [108].

5.1.4 Selection of the regularization parameter

Having computed the LASSO and DS regularization paths for a fMRI time series of a voxel, \mathbf{y} , the sparsity of $\hat{\mathbf{s}}$ depends on the regularization parameter λ as shown in Figure 5.1. As λ is reduced at each iteration, the number of nonzero coefficients in the active set ($\|\mathbf{s}\|_0$) increases such that more degrees of freedom are employed to fit the fMRI signal. Therefore, selecting λ , and in turn the final estimate and signal fit, can be done based on model selection criteria which consider the effective degrees of freedom used to fit the voxel time series.

The effective degrees of freedom of an estimator can be computed as [107]

$$df(\hat{\mathbf{y}}) = \frac{1}{\sigma^2} \sum_{i=0}^{N-1} \text{cov}(\hat{y}_i, y_i). \quad (5.10)$$

where $\hat{\mathbf{y}} = \mathbf{H}\hat{\mathbf{s}}$ is the fitted signal. The linear regression model used in PFM assumes the voxel time series follows a Normal distribution with true mean $E[\mathbf{y}] = \mathbf{H}\mathbf{s}$ and variance σ^2 . Therefore, the covariance between the estimate and the data is given

¹The MATLAB implementation of the algorithm is available at <http://users.ece.gatech.edu/~sasif/homotopy/index.html>

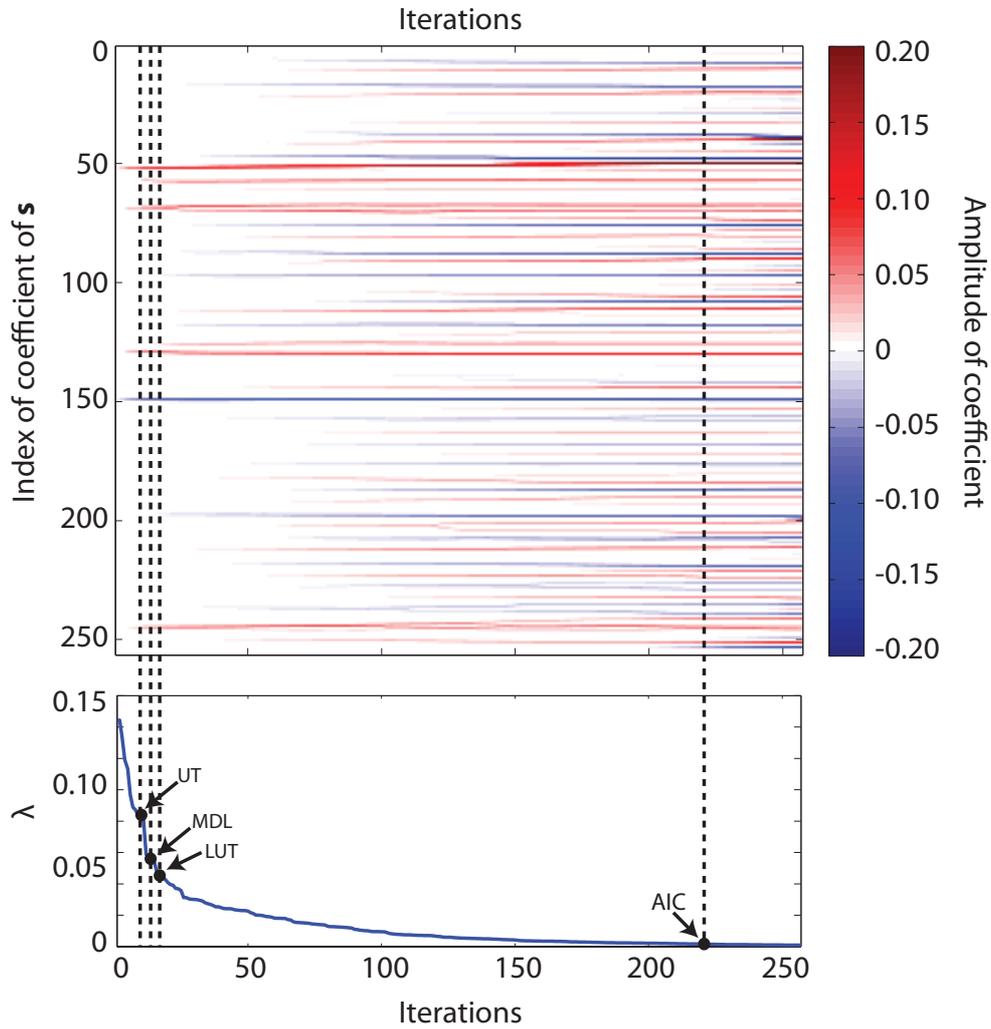


Figure 5.1: Computation of the Regularization path for a simulated fMRI time series (top of Figure 5.3). The top diagram depicts the amplitude of each coefficient in \mathbf{s} (positive amplitude in red, negative amplitude in blue, zero amplitude in white) versus the iteration of the algorithm. The bottom diagram depicts the value of the regularization parameter at each iteration. A coefficient is included or removed from the active set of nonzero coefficients as the regularization parameter λ decreases towards 0 in each iteration of the algorithm. As a consequence, the amplitude of the other coefficients that are at this iteration is also updated accordingly. The first coefficients to be included in the active set correspond to the coefficients that generate the BOLD responses of the simulated voxel time series, as shown in Figure 5.3.

by Stein's unbiased risk estimation (SURE) theory [332]

$$\text{cov}(\hat{y}_i, y_i) = \sigma^2 E \left\{ \frac{\partial \hat{y}_i}{\partial y_i} \right\}, \quad (5.11)$$

providing the mapping $\hat{\mathbf{y}} = f(\mathbf{y})$ is differentiable [107]. For a linear estimator, $\hat{\mathbf{y}}$ can be computed analytically as $\hat{\mathbf{y}} = \mathbf{S}\mathbf{y}$, e.g. $\mathbf{S} = \mathbf{H}(\mathbf{H}^T\mathbf{H} + \lambda\mathbf{I})^{-1}\mathbf{H}^T$ for the ridge regression estimator, such that $\text{cov}(\hat{y}_i, y_i) = \sigma^2 S_{ii}$ and the effective degrees of

freedom for a linear estimator is given by $df(\hat{\mathbf{y}}) = \sum_i S_{ii} = \text{Tr}(\mathbf{S})$ [166].

Nevertheless, obtaining an analytical expression of the effective degrees of freedom for the LASSO and the DS is not so obvious since both estimators are nonlinear subset selection procedures. In general, subset selection algorithms have larger degrees of freedom than the final number of parameters employed to model the signal due to the stochastic search of the best subset among the available coefficients [166]. Interestingly, Zou et al. (2007) demonstrated that the number of nonzero coefficients of $\hat{\mathbf{s}}$ is an exact unbiased estimate of the effective degrees of freedom of the LASSO under SURE's theory providing the design matrix is a full rank matrix, i.e. $\hat{df}(\lambda) = \|\hat{\mathbf{s}}_\lambda\|_0$, where $\hat{\mathbf{s}}_\lambda$ is the estimate for a given λ . Let \mathcal{A}_λ be the active set of nonzero coefficients in $\hat{\mathbf{s}}_\lambda$ for a given λ [384]. According to this result, estimating the number effective degrees of freedom for the LASSO becomes computationally efficient despite being a variable selection procedure.

Nevertheless, the PFM design matrix \mathbf{H} is not orthonormal and thus the previous result cannot be directly applied in SPFM. Here, we proposed to define the following LASSO-based operator

$$\mathbf{S}_{LASSO} = \mathbf{H}_\lambda (\mathbf{H}_\lambda^T \mathbf{H}_\lambda)^{-1} \mathbf{H}_\lambda^T \quad (5.12)$$

where \mathbf{H}_λ is a submatrix of \mathbf{H} with the columns whose coefficients are in the active set \mathcal{A}_λ . Similar to the RR estimator, the effective degrees of freedom used to compute the LASSO estimate can be approximated as

$$\hat{df}_\lambda = \text{Tr}(\mathbf{S}_{LASSO}). \quad (5.13)$$

Numerical simulations demonstrate that the difference between the definition in Eq. (5.13) and the theoretical result in [384] is negligible in event-related fMRI experiments where BOLD signal changes due to neuronal activity are assumed to be non-periodic and sparse in order to maximize the sensitivity of the fMRI experiment [217, 218, 219]. Therefore, it is also particularly well-founded in a single trial fMRI experiment. Furthermore, although the theoretical results in [384] were only developed for the LASSO estimator, we will assume that this approximation can also be employed for fitting the DS since it has been demonstrated that the both solutions become approximately equivalent when \mathbf{s} is likely to be highly sparse [36].

Consequently, given an estimate of the effective degrees of freedom we can formulate adaptive model selection criteria where the optimal value of λ is chosen as

[384]

$$\lambda^* = \arg \min_{\lambda} \|\mathbf{y} - \mathbf{H}\hat{\mathbf{s}}_{\lambda}\|_2^2 + \frac{K}{N}\hat{d}f_{\lambda}. \quad (5.14)$$

In this study we investigated two traditional model selection criteria: the Akaike Information Criterion (AIC) with $K = 2$ [5], and the Minimum Description Length (MDL), which is equivalent to the Bayesian Information Criterion (BIC), with $K = \log N$ [302, 315]. Besides, we also studied to select λ based on the traditional universal threshold (UT) where $\lambda_{UT} = \hat{\sigma}\sqrt{2\log N}$ [101], and the lower universal threshold (LUT) where $\lambda_{LUT} = \hat{\sigma}\sqrt{2\log N - \log(1 + c^2 \log N)}$ ($c = 2$ is used here) [10], and $\hat{\sigma}$ is an estimate of the noise standard deviation. We propose to estimate $\hat{\sigma}$ following the work by Donoho and Johnstone [101]. The fMRI time series is first decomposed with a Daubechies wavelet with 2 vanishing moments. Then, $\hat{\sigma}$ is computed as the median absolute deviance (MAD) of the wavelet coefficients at the finest scale level divided by 0.6745. In many scenarios, the wavelet coefficients at the finest level are essentially pure noise and using the MAD estimate rather than the standard deviation provides control against upward bias due to the presence of some signal at that level [101]. Shrinkage thresholding based on the universal and lower universal thresholds is commonly used in wavelet shrinkage estimators [233].

5.1.5 Debiasing and statistical inference

A debiasing step is usually performed in the active set of coefficients in order to overcome the tendency of the DS and the LASSO to underestimate the true value of the nonzero coefficients (negative bias) [68, 120]. For that, we rely on the DS or the LASSO to define a new linear model which only includes those regressors with nonzero coefficients at the selected solution, \mathbf{H}_{λ} . Futhermore, we also propose to extend the model with a set of additional regressors, \mathbf{X} , which may explain some variability of the time series such as the realignment parameters estimated during rigid-body motion correction. Therefore, we define a debiasing general linear model

$$\mathbf{y} = \mathbf{H}_{deb}\mathbf{s}_{deb} + \epsilon. \quad (5.15)$$

where the debiasing design matrix is defined as $\mathbf{H}_{deb} = [\mathbf{H}_{\lambda} \mathbf{X}]$ and the weights of the regressors for the debiasing model are given by $\mathbf{s}_{deb} = [\mathbf{s}_{\lambda}^T \mathbf{s}_{\mathbf{X}}^T]^T$.

The coefficients of this model are estimated from the fMRI time series via least squares

$$\hat{\mathbf{s}}_{deb} = (\mathbf{H}_{deb}^T \mathbf{H}_{deb})^{-1} \mathbf{H}_{deb}^T \mathbf{y}, \quad (5.16)$$

Algorithm 1: Sparse Paradigm Free Mapping

Parameters: Convolution (Toeplitz) matrix \mathbf{H} according to an HRF model, algorithm: LASSO or DS, stopping factor ν .

For each fMRI voxel time series \mathbf{y} :

Estimate $\hat{\sigma}_{MAD}$ and set $\lambda_{min} = \nu \hat{\sigma}_{MAD}$

Initialization: $\hat{\mathbf{s}} = \mathbf{0}$, $\lambda_{max} = \|\mathbf{H}^T \mathbf{y}\|_\infty$

while $\lambda^{(n)} < \lambda_{min}$ **do**

- Obtain the coefficient path of LASSO or DS solutions with the PD pursuit homotopy algorithm [15] until $\lambda < \lambda_{min}$. Solutions are computed at critical values of λ , where the active set \mathcal{A}_λ is updated;
- Estimate $\hat{df}(\lambda)$ with (5.13) and compute $\|\mathbf{y} - \mathbf{H}\mathbf{s}_\lambda\|_2^2$ for each solution;
- Select the optimal λ with MDL, AIC, UT or LUT, and define the final active set of coefficients \mathcal{A}_λ ;

Define the debiasing linear model, \mathbf{H}_{deb} , and regress it out against the voxel time series via Least Squares;

Compute t - and F -statistics of the regressors included in \mathbf{H}_{deb} and set t -values of the zero coefficients equal to zero;

whereas the coefficients not included in \mathcal{A}_λ are set to zero.

Assuming that the noise follows a Gaussian distribution and conditional on the debiasing model, summary statistics assessing the significance of the different regressors in \mathbf{H}_λ and \mathbf{X} can be derived with standard t - and F -tests. Here, we assume that the zero coefficients in the LASSO or DS solution, i.e. those not included in \mathcal{A}_λ , are not statistically significant and have a t -value equal to zero [339].

Pseudocode Algorithm

A pseudocode implementation of Sparse Paradigm Free Mapping is given in Algorithm 1. The regularization path is computed for each fMRI voxel time series and the final solution is then selected depending on the method to choose λ . In practice, the computational cost of the procedure can be reduced with some heuristic rules. For instance, the homotopy algorithm can stop when λ is below a certain λ_{min} , e.g. a factor ν of the estimated noise standard deviation $\hat{\sigma}_{MAD}$, or when the number of coefficients included in the active set is higher than a sparsity threshold, e.g. the number of nonzero coefficients cannot be larger than half of the length of the fMRI time series N .

5.2 Methods

Simulated fMRI time series and experimental fMRI data were used to evaluate the performance of the LASSO and the DS algorithms in SPM, and analyze

the differences between using AIC, MDL, UT and LUT to select the regularization parameter λ . In all experiments, the HRF convolution matrix \mathbf{H} used in the estimation was defined from the two gamma-variate HRF with standard SPM (FIL/UCL) parameters (canonical HRF) [130] sampled at the corresponding TR.

5.2.1 Simulated fMRI data

We generated $N_b = 1000$ fMRI time series for each simulation scenario which was characterized by the number of events of the simulated fMRI series, the HRF shape and the tSNR. The BOLD component of the fMRI time series was created as the convolution of a stimulus time series, $s(t)$, with an HRF, $\tilde{h}(t)$. Six different HRF shapes for $\tilde{h}(t)$ were simulated based on the canonical HRF with varying time-to-peak from 3 s to 8 s, as shown in Figure 5.2, whereas the standard shape of the canonical HRF used to define the matrix \mathbf{H} has a time-to-peak of 5 s. The mismatch between the simulated HRF and the model HRF allowed us to examine the effect of model deviations in the detection and estimation performance of the technique. Both $\tilde{h}(t)$ and $s(t)$ were initially created with a temporal resolution of 10 points per TR in order to approximate them as continuous signals. The duration of the simulated fMRI time series was 256 s. For the stimulus signal, zero (no activations) to ten events of duration 2 s and constant amplitude but with random polarity were generated, which allowed us to investigate how the different criteria perform with a variable level of sparsity of the stimulus signal. The onset of the events was randomly sampled from a uniform distribution across the duration of the time series at the pseudo continuous temporal resolution of the signals (TR/10). After convolution, the simulated BOLD signal was downsampled to TR of 2 s, i.e. $N = 128$ time points.

Two noise models were investigated in the simulations: additive white gaussian noise (AWGN), and a more realistic scenario where the noise consisted of AWGN plus a sinusoidal term, aiming to represent cardiac and respiratory fluctuations in the fMRI signal [144]. The sinusoidal term was generated as:

$$\sum_{i=1}^{N_h} \frac{1}{2^{i-1}} \left(\sin(2\pi f_{r,i}t + \phi_{r,i}) + \sin(2\pi f_{c,i}t + \phi_{c,i}) \right). \quad (5.17)$$

with 4^{th} -order harmonics ($N_h = 4$) per component. The frequency of each harmonic was randomly generated with a Normal distribution with mean $if.$ and variance 0.04, i.e. $f_{.,i} \sim N(if., 0.04)$, and the fundamental frequencies were $f_r = 0.3$ Hz for

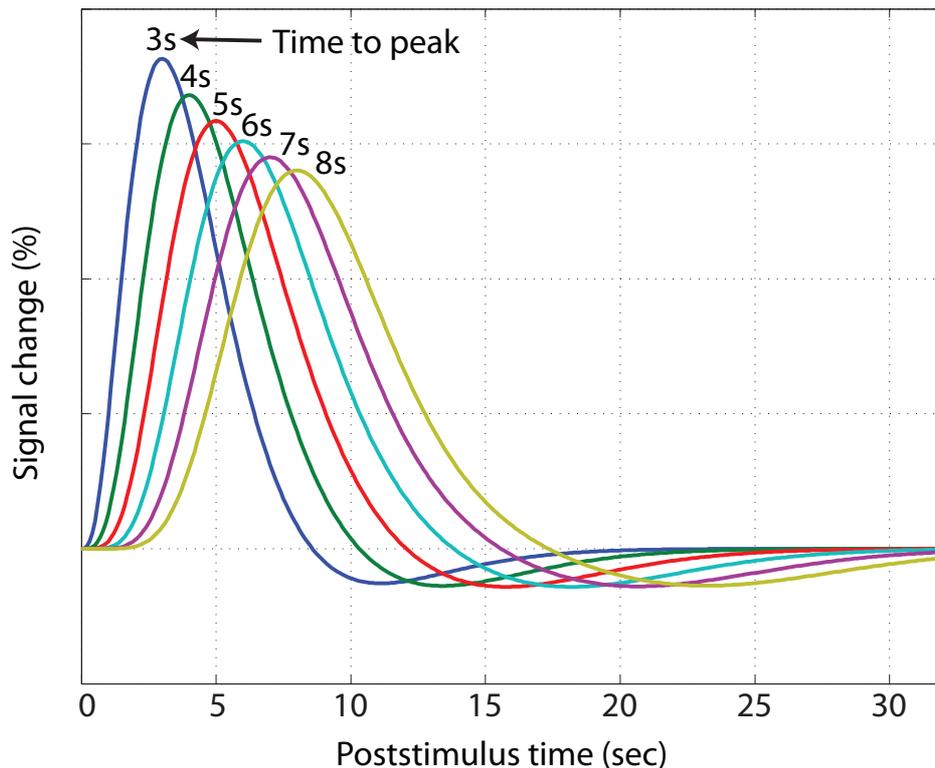


Figure 5.2: Haemodynamic response functions (HRF) used in the simulations. These models of the HRF are based on the canonical HRF with standard SPM parameters but varying the time-to-peak between 3 s and 8 s. The standard value of time-to-peak is 5 s.

the respiratory components [38] and $f_c = 1.1$ Hz for the cardiac components [320]. A power-of-2 decrease in the amplitude of harmonics was simulated following the results in [144]. Temporal signal-to-noise ratios (tSNR) between 30 and 80 were studied which assume a signal change of 6% at 7T [350] and correspond to contrast-to-noise ratios (CNR) between 1.8 and 4.8 [200, 343]. Since the physiological noise is proportional to the BOLD signal change [200], we also modelled a variable ratio between the sinusoidal and AWGN components of the noise, representing physiological (σ_P) and thermal (σ_0) noise, respectively. The model was $\sigma_P/\sigma_0 = a \cdot SNR^b + c$, where $a = 5.01 \cdot 10^{-6}$, $b = 2.81$ and $c = 0.397$. This model was extracted from the experimental measures of the physiological to thermal noise ratio presented in Table 3 in Triantafyllou et al. (2005) [343] at 7T. The noise component was added to the simulated signal to create a fMRI time series, which was analyzed with the SPM method for each possible combination of algorithm (DS and LASSO) and criteria (UT, MDL, LUT and AIC).

Two measures were used to summarize the results of the analysis. We computed the mean square error between the fitted ($\mathbf{H}\hat{\mathbf{s}}$) and the simulated haemodynamic

signals ($\tilde{\mathbf{H}}\mathbf{s}$), i.e.

$$\text{MSE}_s = \frac{1}{N_b} \sum_{i=1}^{N_b} \|\mathbf{H}\hat{\mathbf{s}}_i - \tilde{\mathbf{H}}\mathbf{s}_i\|_2^2. \quad (5.18)$$

In addition, defining a false positive (FP) event (Type-I error) when a nonzero coefficient in the estimated signal did not correspond to an 'ON' time point in the simulated stimulus signal, and a false negative (FN) event (Type-II error) when an 'ON' event in the simulated \mathbf{s} did not correspond with a nonzero coefficient in $\hat{\mathbf{s}}$, we measured the sensitivity and specificity of the technique to detect the simulated activation events. The specificity and sensitivity ratios were defined as

$$\text{Specificity} = \frac{\text{Number of TNs}}{\text{Number of TNs} + \text{Number of FPs}} = 1 - \text{FPR}, \quad (5.19)$$

and

$$\text{Sensitivity} = \frac{\text{Number of TPs}}{\text{Number of TPs} + \text{Number of FNs}} = 1 - \text{FNR}, \quad (5.20)$$

where FPR and FNR denotes the False Positive Rate and False Negative Rate, respectively. Therefore, both rates represent temporal sensitivity and specificity, rather than spatial as is usually done in fMRI analysis. The tradeoff between sensitivity and specificity was summarized in terms of Receiver Operating Characteristic (ROC) curves [189].

The specificity and sensitivity measures allow us to evaluate the ability of the method to detect the temporal location of the nonzero coefficients of the stimulus function that was simulated to create the BOLD signal. On the other hand, the measures of MSE examine the accuracy of the computed estimates of the BOLD signal [204] and allow us to evaluate the robustness of the method to recover the BOLD component of the signal in case of uncorrect nonzero coefficients in the estimated stimulus signal (false negatives and false positives). Alternatively, one could consider to assess the quality of the estimation in terms of the variance explained by the estimate or the coefficient of determination R^2 [297].

For comparison the time series were also analyzed with a GLM comprising one regressor per activation event, which was created as the convolution of the Canonical HRF with the simulated stimulus function of the event. This is an ideal non paradigm free scenario where the onsets of the activations are known a priori and therefore it is optimal assuming no HRF mismatch exists with respect to the model.

5.2.2 Experimental fMRI data

SPFM was evaluated with two different experimental paradigms. We analyzed the datasets acquired with TR 2 s of the visuomotor paradigm with visually cued and self paced single trial finger tapping events described in Chapter 4 in order to examine the differences between both PFM methods. As noted in that chapter, surface electromyography (EMG) signals were recorded to capture the timing of the self paced finger tapping trials. Second, we analyzed datasets acquired in one patient with idiopathic generalized epilepsy in order to determine the ability of SPFM to detect interictal epileptiform discharges (IED) without relying on MR-compatible EEG-determined timing information about the onsets of the seizures, and contrast the activations detected with SPFM with those detected with an EEG-based GLM analysis.

Visuomotor paradigm: Acquisition and Preprocessing

We refer the reader to the methods section in Chapter 4 for information about the MRI acquisition, description of the experimental paradigm, preprocessing steps (motion correction, detrending, amplitude normalization to percentage signal change) and recording of the EMG signals. Similar to Chapter 4, we computed Spearman's non-parametric rank correlation coefficient [171] between the Activation Time Series (ATS) obtained with SPFM and EMG signals. The EMG signals were decimated to the same number of points as the ATS and thresholded at an amplitude z -score of 4, i.e. 4 times the standard deviation from the EMG time course mean, in order to only consider significant movements.

Epilepsy data: fMRI/EEG Acquisition and Preprocessing

Eight runs of fMRI data were acquired in one session in a patient with idiopathic generalized epilepsy and a WM lesion in the dorsal part of the right hemisphere apparent on the EPI T_2^* -weighted functional images (see Figure 5.10). The scanning session was conducted at the University of Birmingham using a 3T Philips Achieva scanner (data courtesy of Dr. Andrew Bagshaw). Each run lasted for 375 s, recording 125 single-shot gradient echo EPI scans (TE 35 ms, SENSE factor 2, 2.5 mm isotropic voxels) sampled at TR 3 s with 49 slices. The first 5 scans were discarded to allow steady-state magnetization to be achieved resulting in a time series with 120 scans. Slices were oriented parallel to the anterior commissure - posterior commissure

(AC-PC) axis of the patient’s brain and were positioned superior to the pons and extending to the whole brain.

Prior to the analysis with GLM or SPFM, datasets from each run were motion corrected with rigid body registration, and detrended by regressing out the voxel time series with the first 4th-order Legendre polynomials and the sine and cosine signals with one cycle over the scan duration. The voxel time series were normalized in order to represent percentage signal change based on the mean amplitude of the time series. The data were also spatially smoothed with a Gaussian filter with 4 mm FWHM. These steps were performed using AFNI (NIMH/NIH) [82].

EEG data was simultaneously acquired (EEG-data acquisition and analysis was done by Dr. Andrew Bagshaw) from 64 electrodes using an MR compatible EEG system (BrainProducts, Munich, Germany), sampled at 5000 Hz and filtered with a band pass filter with cutoff frequencies between 0.1 Hz and 250 Hz. The EEG data acquisition setup clock was synchronized with the MRI scanner clock facilitating the location of the IED onsets in the fMRI data. Raw EEG data were partitioned into the fMRI runs and preprocessed for gradient artefact removal with BrainVision software [6] and ballistocardiographic artefact removal using the Optimal Basis Set method [264]. Interictal discharges were identified in the cleaned EEG data.

Sparse PFM analysis

Datasets of both experiments were analyzed with the SPFM method implemented with in-house code written in Matlab. Based on the results on simulated data, datasets were only analyzed with the DS estimator considering the four criteria MDL, AIC, UT, LUT. To define the debiasing matrix in 5.15, we also included the 6 translation and rotation realignment parameters estimated with rigid body registration as additional nuisance regressors in the matrix \mathbf{X} (see Eq. (5.15)). After fitting the debiasing model, t -statistics of HRF-based coefficients were computed in order to assess the significance of detected activations. Note that the part of the debiasing model describing the BOLD events \mathbf{H}_{deb} is different for each voxel according to the voxelwise estimate computed with the DS or the LASSO. The use of different models for each voxel has been previously proposed to overcome suboptimal model specification that can occur if a unique, usually over-simplified model which may not capture the voxelwise variability of the fMRI data is used [230, 297]. Nevertheless, direct visualization and comparison of the statistics with voxelwise variability in the degrees of freedom of the statistics is inappropriate. As a solution, statistics

were standardized to z -scores considering the degrees of freedom of each voxel and p -values were computed accordingly. When the statistics were corrected for multiple hypothesis testing, we used a modified FDR procedure that first estimates the number of true negatives (m_0) and then adjusts the FDR corrected Q -values downward according to $Q^* = Q \cdot m_0/m$, where m is the number of coefficients to test excluding those coefficients with $z = 0$ [33]. This procedure is the default approach for FDR control in the AFNI 3dFDR function.

We relied upon the concepts of Activation Time Series (ATS) and Activation Events presented in Chapter 4 in order to explore the SPFM results in time without using information from the EMG in the visuomotor paradigm or the EEG in the epileptic study. Total, positive and negative ATS were computed for two thresholding cases: no-thresholding and FDR-corrected $Q < 0.05$. The ATS with no thresholding counts all the voxels whose z -score at each time point is nonzero (equal to nonzero estimated coefficient) and therefore it aims to mimic the scenario of the simulations.

In practice, a completely paradigm free approach would require one to study all the activation events detected in the ATS. In this study, we only studied those activation events corresponding to significant hand movements in the EMG signals for the visuomotor paradigm and the IEDs detected in the EEG for the epileptic datasets. In addition, we also studied those activation events showing a peak in the thresholded ATS. To ease the identification of the areas showing activations in each activation event, we computed an SPFM Activation Map for each event with the maximum absolute z -score of the time points which were marked in the thresholded ATS. However, the SPFM Activation Maps plotted the z -scores without applying any amplitude threshold in order to maximize sensitivity, but a minimum cluster size of 2 voxels was set to reduce isolated activations.

GLM analysis

The datasets were also analyzed with a standard GLM where the regressors describing the experimental effects were created by the convolution of the event related stimuli with the SPM canonical HRF and its temporal and dispersion derivatives (3 regressors per event), and the 6 realignment parameters as nuisance covariates [130]. The events in the stimulus time series were modelled as delta functions at the times of onset of finger tapping shown in the EMG recordings, and the times of IEDs marked in the EEG. GLM analysis was performed with the AFNI 3dDeconvolve

function. This function does not consider noise serial correlations and was used to make the analysis similar to the SPM method which assumes no correlation model for the noise. Clusters of activation needed to have a minimum of two contiguous voxels with significant F -statistic of the regressors describing the event ($Q < 0.05$, FDR corrected). However, GLM maps plotted the t -statistic corresponding to the canonical HRF regressor of the event in order to differentiate between positive and negative BOLD signal changes.

In order to examine the degree of concordance between the GLM and SPM maps, we calculated the number of overlapping voxels in both maps for each of the tapping events ($\text{SPM} \cap \text{GLM}$), and the percentage of overlapping voxels relative to the number of detected voxels with SPM ($\% \text{SPM in GLM} = \text{SPM} \cap \text{GLM} / \text{SPM}$).

Detection of sudden head motion

In order to examine whether the detected activations may originate from severe signal changes due to head motion, we estimated the magnitude of the net displacement vector, d , from the translation parameters estimated during motion correction. The absolute scan-to-scan displacement was then computed as the absolute derivative of d , $|d'|$. Time points with significant head motion were declared when $|d'| > 0.5$ mm/scan and marked in the ATS plot [209].

5.3 Results

5.3.1 Simulated data

The results of the simulations illustrated the potential of using sparse estimation techniques to detect event related single trial BOLD responses in a complete paradigm free framework. Figure 5.3 illustrates how SPM operates and shows the differences between the four methods to select λ investigated in this work. In this example, the simulated fMRI time series included 5 activation events of duration 2 s. After computing the regularization path, the solutions computed with the use of the UT and MDL criteria correctly detect the 5 simulated events, whereas setting λ to the LUT causes a false positive detection at approximately 200 s. The figure also demonstrates that AIC tends to select a very low λ resulting in data overfitting.

Figure 5.4 depicts the results of the ROC analysis when the DS (left figures) and the LASSO (right figures) is employed in SPM for the scenarios with AWGN noise (top figures) and AWGN plus sinusoidal trends (bottom figures). The pro-

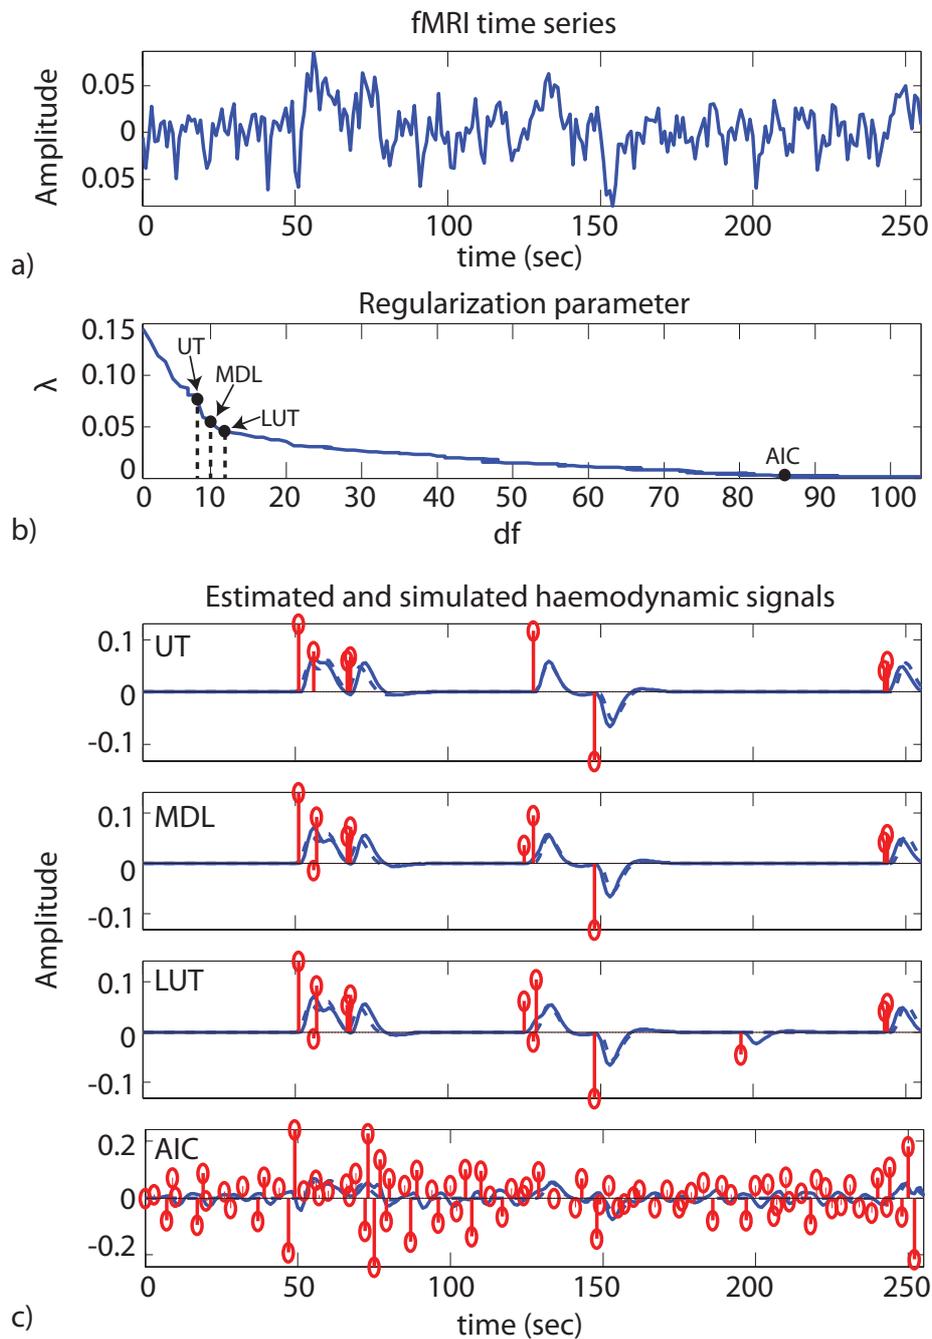


Figure 5.3: (a) For each voxel fMRI time series, (b) the regularization path is obtained and the effective degrees of freedom is computed for each possible solution. (c) Model selection and thresholding criteria are used to select the regularization parameter, and in doing so the final estimate of s (red) and signal fit (blue) are chosen. In this figure, the fMRI time series (128 time points with TR 2 s) includes 5 activation events of duration 2 s and simulated tSNR is 50 with AWGN and sinusoidal noise. The HRF used to generate the fMRI signal is different from the HRF used by the algorithm (HRF mismatch). It can be seen that the UT ($\lambda=0.080$) and MDL ($\lambda=0.054$) criteria correctly detected the 5 activation events, but the UT selected less nonzeros (2 active coefficients per event) than MDL. The LUT ($\lambda=0.045$) criteria incorrectly detects an event around 200 s, and the AIC solution ($\lambda=0.001$) completely fails to estimate the stimulus signal and overfits the fMRI time series.

posed method achieves false positive rates (FPR) smaller than 5-7% for the MDL, LUT and UT criteria when the simulated HRF shape is equal to the HRF of the model, i.e. in a no HRF-mismatch scenario. As the fMRI signal becomes more noisy (tSNR decreases from 80 to 30), the sensitivity of the method decreases but importantly the specificity is not reduced. The specificity of the method diminishes drastically with AIC with little increased sensitivity. It can be observed that the DS slightly outperforms the LASSO in terms of specificity for the MDL, LUT and UT, whereas the advantage of the DS becomes more significant when using AIC which select less sparse solutions. It is shown that the technique exhibits slightly improved performance (1-2% specificity improvement) when the noise includes AWGN and sinusoidal components (bottom figures) rather than AWGN (top figures). The results clearly demonstrate that the sensitivity of the method diminishes when the simulated HRF substantially differs from the one modelled, i.e. in an HRF-mismatch regime. In these figures, the curves at the bottom of each figure correspond to simulated HRFs with time-to-peak of 8 s whereas the modelled HRF has a time-to-peak of 5 s. In a mismatch scenario, sensitivity values larger than 0.4 are not feasible with any method for this particular simulation. Nevertheless, it can be observed that the method is still able to operate at FPR values lower than 5 % with MDL, LUT and UT at low tSNR conditions, suggesting that the technique increases the threshold for detection at low tSNR in order to avoid detecting false positives. Comparing the ROC curves for varying number of simulated events (2 events of duration 2 s shown with \circ markers, 6 events shown with \times markers, and 10 events shown with \square markers), it can be seen that for all the combinations of algorithm, criteria and type of noise that the technique operates at better specificity-sensitivity points when the stimulus signal has a very small number of events, i.e. it is very sparse.

The mean square error between the simulated BOLD signal and its estimate (MSE_s) is plotted as a function of the number of simulated events in Figure 5.5 and as a function of tSNR in Figure 5.6. Similar to the ROC analysis, both figures demonstrate no significant difference between both estimators (solid lines for LASSO and dash lines for DS) except with AIC, which tend to overfit the time series by detecting a large number of activations. More accurate estimates are obtained in the case of AWGN plus sinusoidal trends. Importantly, no substantial reduction in the accuracy of the estimates is observed when there is a mismatch in the HRF model (curves with \square markers correspond to simulated HRF with time to peak of 8 s, whereas curves without markers correspond to perfect knowledge of the HRF

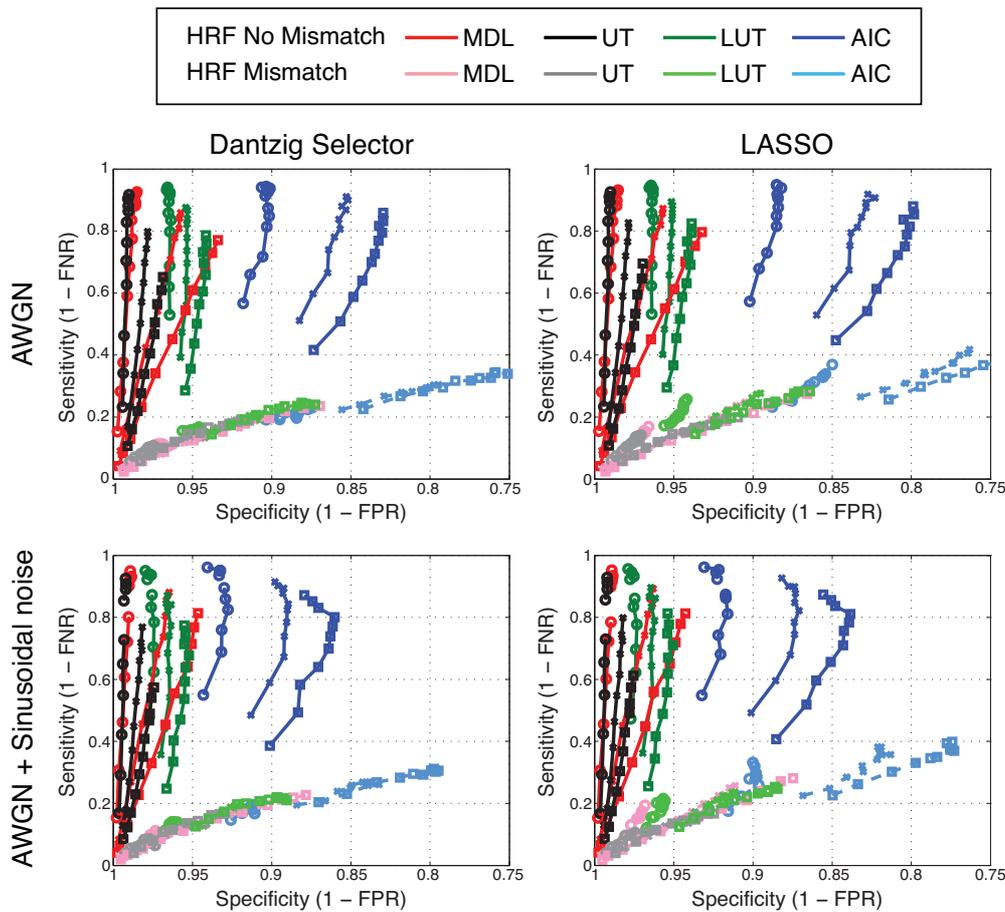


Figure 5.4: Receiver Operating Curves (Sensitivity vs. Specificity) of SPFM. The model HRF is the Canonical HRF with standard SPM parameters. No mismatch scenario: simulated HRF is equal to the model HRF with time-to-peak of 5 s; HRF Mismatch scenario: simulated HRF has a time to peak of 8 s whereas the model HRF has a time-to-peak of 5 s (see Figure 5.2). Line markers indicate number of activation events: 2 (\circ), 6 (\times) and 10 (\square). Points with decreasing sensitivity in each curve correspond to tSNR decreasing from 80 to 30.

with time to peak of 5 s). This robust characteristic of the technique in terms of accuracy of estimate of the simulated BOLD response contrasts with the one observed in the sensitivity and specificity analysis, suggesting that SPFM reconstructs the stimulus signal accurately even though the HRF used to define the matrix \mathbf{H} does not properly model the HRF. Comparing the right and left plots of Figure 5.5 and the plots in Figure 5.6 across rows, one can observe that the estimation of the stimulus signal improves as tSNR increases but the relative improvement depends on the number of activation events. UT and MDL produce more accurate estimates with highly sparse stimulus signals. Decisively, only MDL, and UT to a lesser extent, detect no responses in the absence of activations. Nevertheless, the large regularization parameters λ selected with MDL and UT makes the accuracy of the

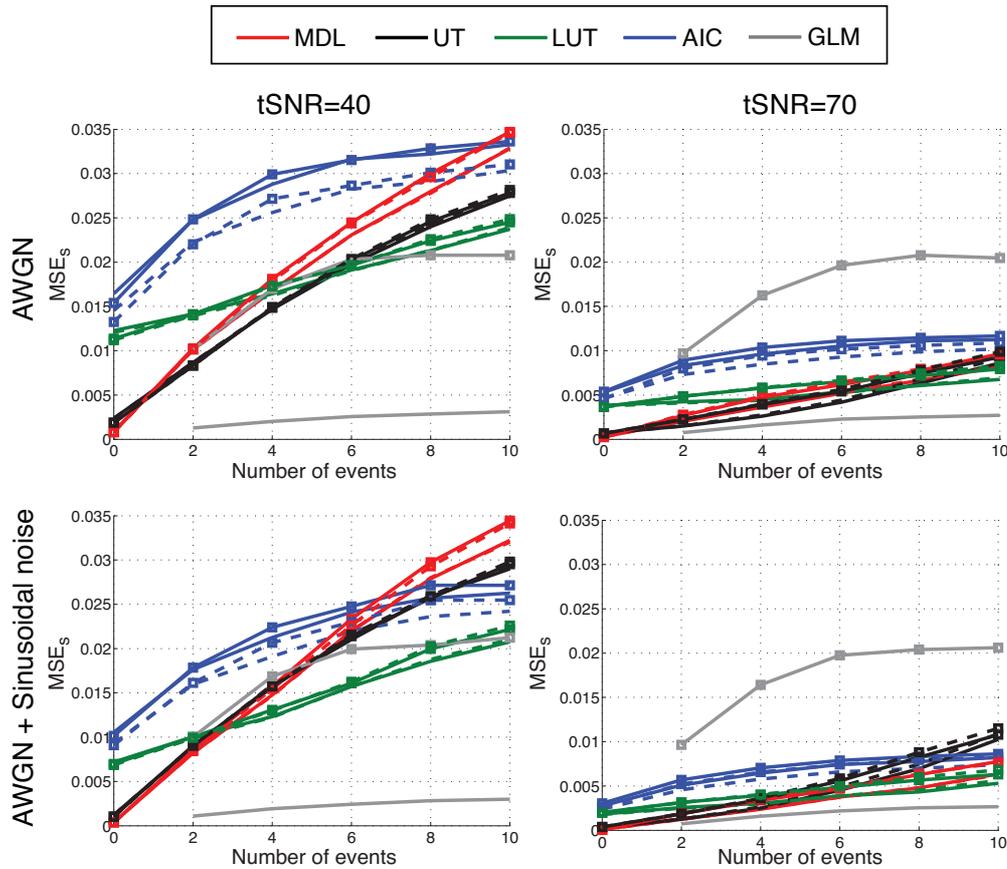


Figure 5.5: Mean Square Error between the simulated and the estimated BOLD signal (MSE_s) as a function of the number of simulated activation events. LASSO: solid line; DS: Dash line. Curves with square markers (\square) correspond to HRF Mismatch (HRF time to peak of 8s), whereas curves without markers correspond to no HRF mismatch.

estimation degrade in less sparse scenarios where some events may be undetected due to excessive shrinkage of the coefficients. In contrast, LUT offers better results in time series with numerous events but at the cost of detecting responses in the absence of events, i.e. detecting false positives. The GLM curves in the case of no HRF mismatch serve in this case as asymptotic limits to the L_1 -norm estimators. The MDL and UT criteria operate close to this limit for high tSNR and very sparse signals. However, the accuracy of the estimates obtained with the GLM approach considerably deteriorates in the case of HRF discrepancy, especially in a high tSNR scenario when the noise includes sinusoidal trends.

5.3.2 Experimental data

The SPFM technique using the MDL criterion to select the regularization parameter demonstrated the best tradeoff between sensitivity and specificity to detect

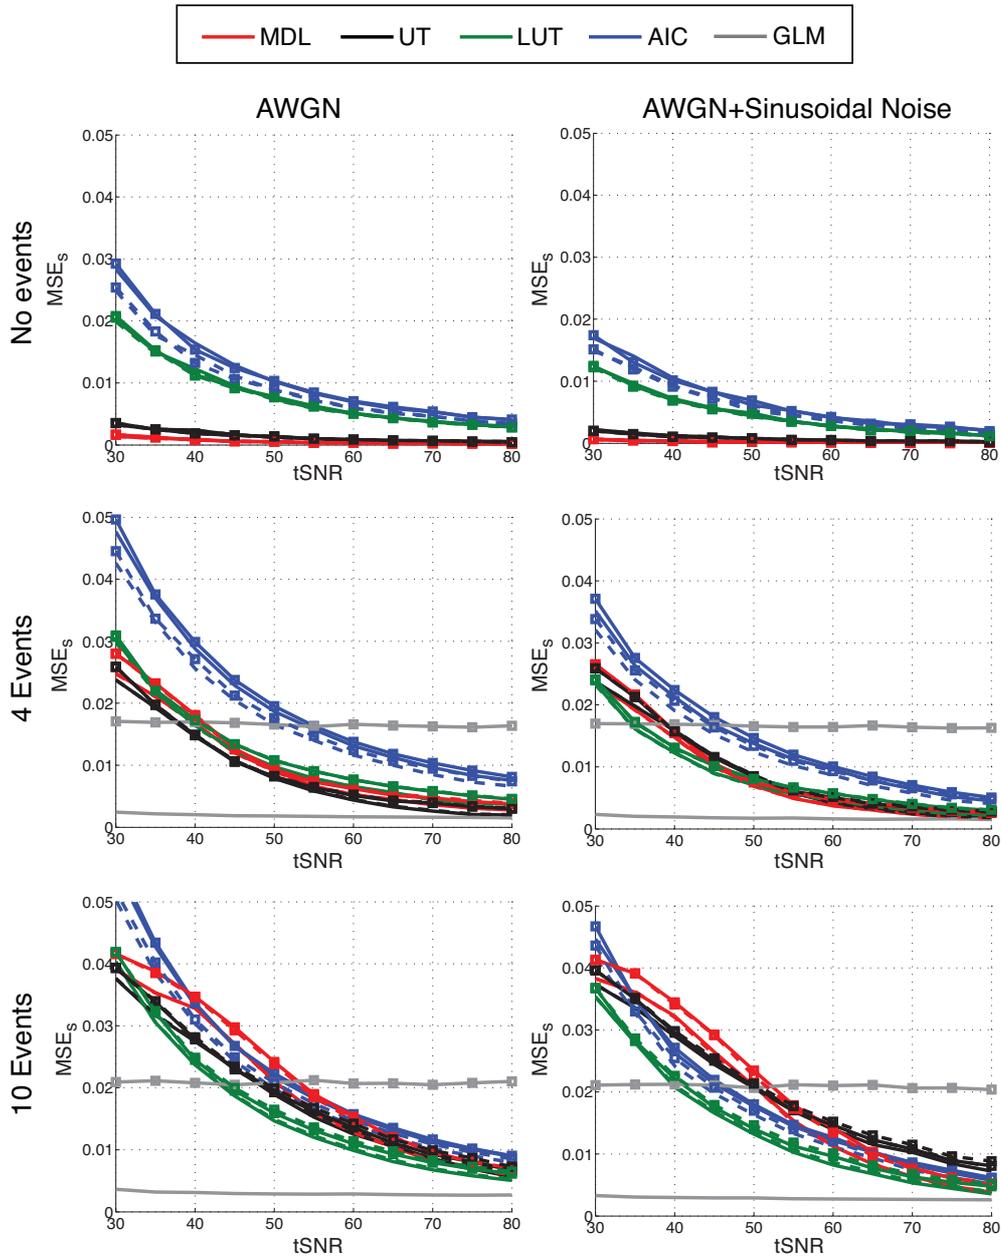


Figure 5.6: Mean Square Error between the simulated and the estimated BOLD signal (MSE_s) as a function of tSNR. LASSO: solid line; DS: Dash line. Curves with square markers (\square) correspond to HRF Mismatch (HRF time to peak of 8s), whereas curves without markers correspond to no HRF mismatch.

single trial event related BOLD responses in both 3T and 7T experimental scenarios (Figures 5.7 and 5.9).

Visuomotor paradigm

Figure 5.7 plots the non-thresholded ATS (ATS_{NoTH}) and thresholded ATS (ATS_{TH}) computed using the DS and the MDL for all the visuomotor datasets,

	Left Extensor		Right Extensor		Right Flexor	
	NoTH	TH	NoTH	TH	NoTH	TH
Subject A	0.210*	0.222*	0.271*	0.267*	0.366*	0.287*
Subject B	0.042	0.052	0.179*	0.177*	0.287*	0.216*
Subject C	0.193*	0.207*	0.195*	0.065	0.175*	0.074
Subject D	0.310*	0.347*	0.293*	0.288*	0.228*	0.254*
Subject F	-0.076	-0.071	0.062	0.051	0.174*	0.176*

Table 5.1: Spearman’s rank correlation coefficients between the decimated EMG signals and the ATS computed with and without FDR statistical thresholding. The asterisks indicate statistically significant correlation coefficient ($p < 0.001$).

along with the corresponding EMG left extensor (LE) and right flexor (RF) signals. No significant head movements ($|d'| > 0.05$) were observed in any datasets. It can be observed that all single trial finger tapping responses are detected in the non-thresholded ATS. Nevertheless, the sensitivity of the technique to detect single trial BOLD signal changes is observed to considerably decrease when the nonzero z -scores were statistically thresholded at $Q < 0.05$, FDR-corrected. The non-thresholded ATS computed for the remaining criteria (UT, LUT and AIC) showed a higher number of active voxels than MDL at any time point. Notably, no activation events were found in the ATS of the AIC and LUT after thresholding, and the ATS computed with UT indicated activation events with less than 25 active voxels at the same times found with MDL (either related to the finger tapping task or at rest (Figure 5.7), suggesting that the nonzero coefficients estimated with these criteria presented lower statistical significance due reduced degrees of freedom of the statistics.

Consequently, Spearman’s non-parametric rank correlation coefficients between the non-thresholded and thresholded ATS and the EMG signals were only computed for the MDL criteria. Table 5.1 illustrates that the non-thresholded ATS were significantly correlated with the EMG signals ($p < 0.001$) except when no significant hand movements were recorded (left extensor (LE) and right extensor (RE) of Subject E, and LE of Subject B). Additionally, no significant correlation was found between the thresholded ATS and right extensor (RE) and right flexor (RF) of Subject C where the ATS did not show activations at the time of the visually-cued finger tapping events.

Figure 5.8 depicts the SPFM maps (using DS and MDL criteria) and the GLM maps ($Q < 0.05$, FDR-corrected) for the finger tapping events of subject A. The SPFM maps were created without statistical thresholding, i.e. each map shows maximum z -score of the time points detecting the event in the ATS. It can be seen

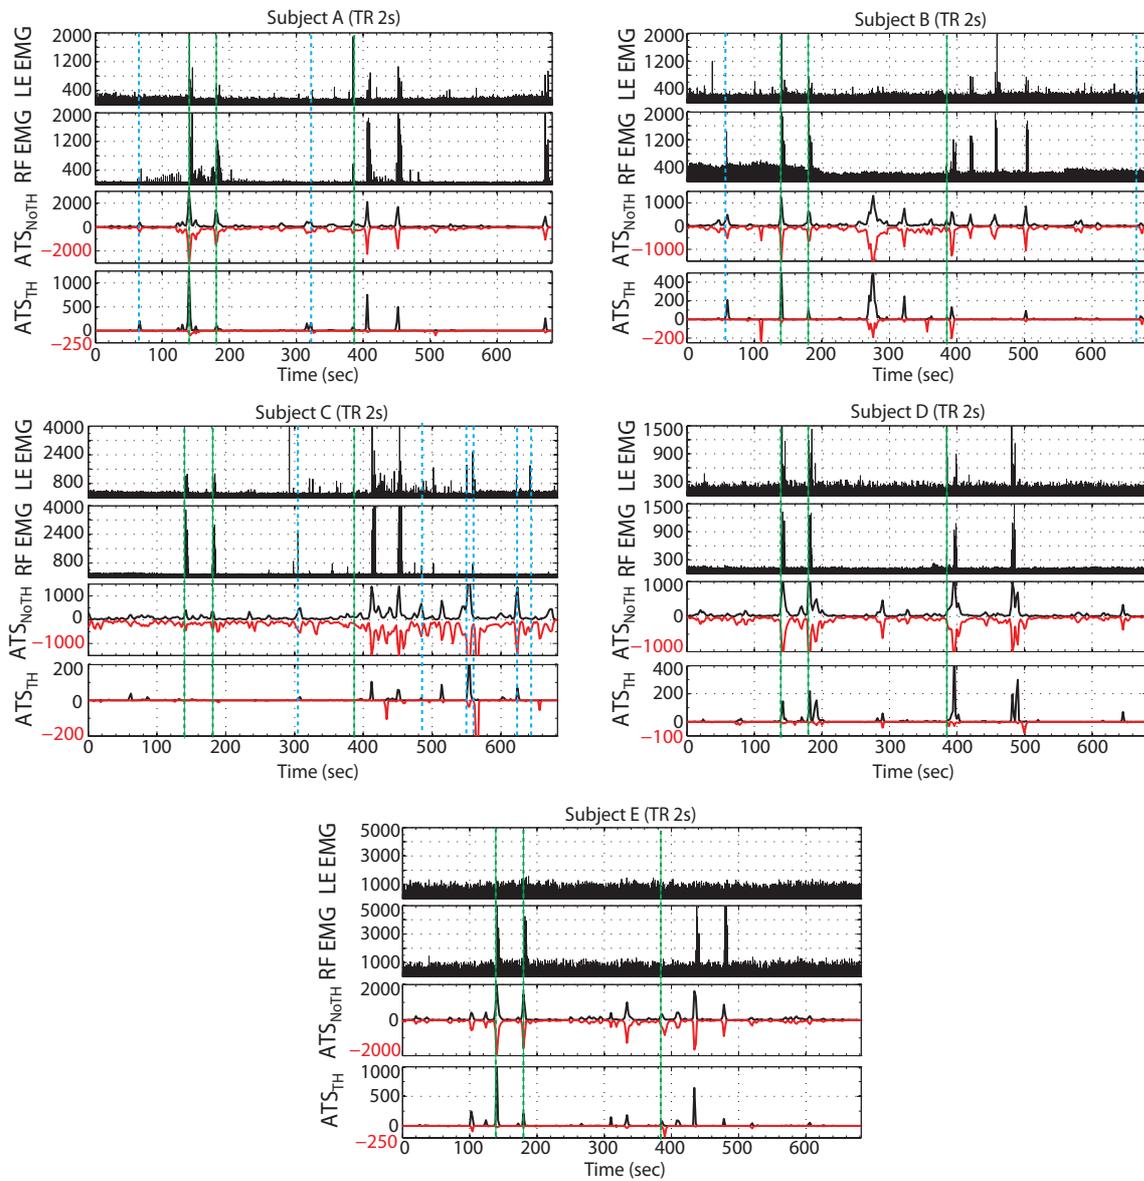


Figure 5.7: Activation Time Series (ATS) computed with SPFM using the DS algorithm with the MDL selection criteria. Each plot shows the EMG time series corresponding to the left extensor (LE) (first row) and right flexor (second row), the fMRI Activation Time Series before thresholding (ATS_{NoTH} , third row) and after thresholding the z -scores to FDR-corrected $Q < 0.05$ (ATS_{TH} , fourth row). In order to facilitate the interpretation of the result, the ATS differentiates between positive activations (black, positive y-axis) and negative activations (red, negative y-axis). The green dashed lines indicate the times of visual cues for finger tapping or start of self-paced part of the run. As illustrated, all main finger tapping events are detected without thresholding. Nevertheless, the visually cued taps of subject C and the third and fourth self-paced taps of subject B are not revealed in the thresholded ATS. In addition, SPFM was able to detect some spiking movement recorded in the EMG as marked with blue dashed lines.

that the clusters of activations shown in the SPFM maps closely matched those revealed in the GLM maps, suggesting high spatial specificity of the activations detected with SPFM relative to the activations detected with GLM. Equivalent results were obtained for the remaining datasets. Averaging across all finger tapping events and all subjects, we observed that 75.6% in mean (or 81.9 % in median) of the voxels identified as active in the SPFM maps were also present in the GLM maps.

The activations related to the finger tapping events were located in areas of visuomotor processing including: the supplementary motor areas (SMA), premotor (PM), primary motor areas (M1), primary somatosensory areas (S1), superior parietal cortex and visual cortex (Figure 5.8). In addition, sparse activations were observed at rest in cortical areas characteristically included in resting state networks, such as the precuneus, posterior cingulate, lateral and superior parietal cortex and superior frontal gyrus, or unilateral or bilateral sensorimotor areas at time points with concordant spikes in the EMG signals (marked with blue lines in Figure 5.7).

Epileptic data

Figure 5.9 depicts the ATS before thresholding (top) and after thresholding ($Q < 0.05$, FDR-corrected) (bottom) for the MDL (solid line) and the UT criteria (dotted line). The interictal epileptiform discharges (IED) revealed by the EEG are marked with blue lines, and the times of significant head movement ($|d'| > 0.5$ mm/scan) are marked with green lines. Only 5 of the 8 runs showed epileptic-related spikes in the EEG data with 8 IEDs identified in total. The ATS computed with SPFM were able to capture 5 after thresholding the z -scores to FDR-corrected $Q < 0.05$, but importantly the non-thresholded ATS revealed 7 out of the 8 IEDs. In general, the SPFM showed good concordance with the EEG-based GLM maps for the single IED events, especially for those IEDs which were apparent in the ATS after thresholding.

Case Reports

Run 1

One IED was identified in the EEG at 3.5 s. This IED was identified at both non-thresholded and thresholded ATS. The corresponding SPFM maps showed activation in the frontal lobe (inferior, middle and superior frontal cortex), insular gyrus, caudate nucleus, anterior middle temporal gyrus, posterior inferior temporal gyrus, occipital lobe, and limbic structures in the medial part of the superior temporal lobe (entorhinal cortex). In addition, negative BOLD changes were located in the posterior inferior and middle temporal gyrus, posterior cingulate gyrus and

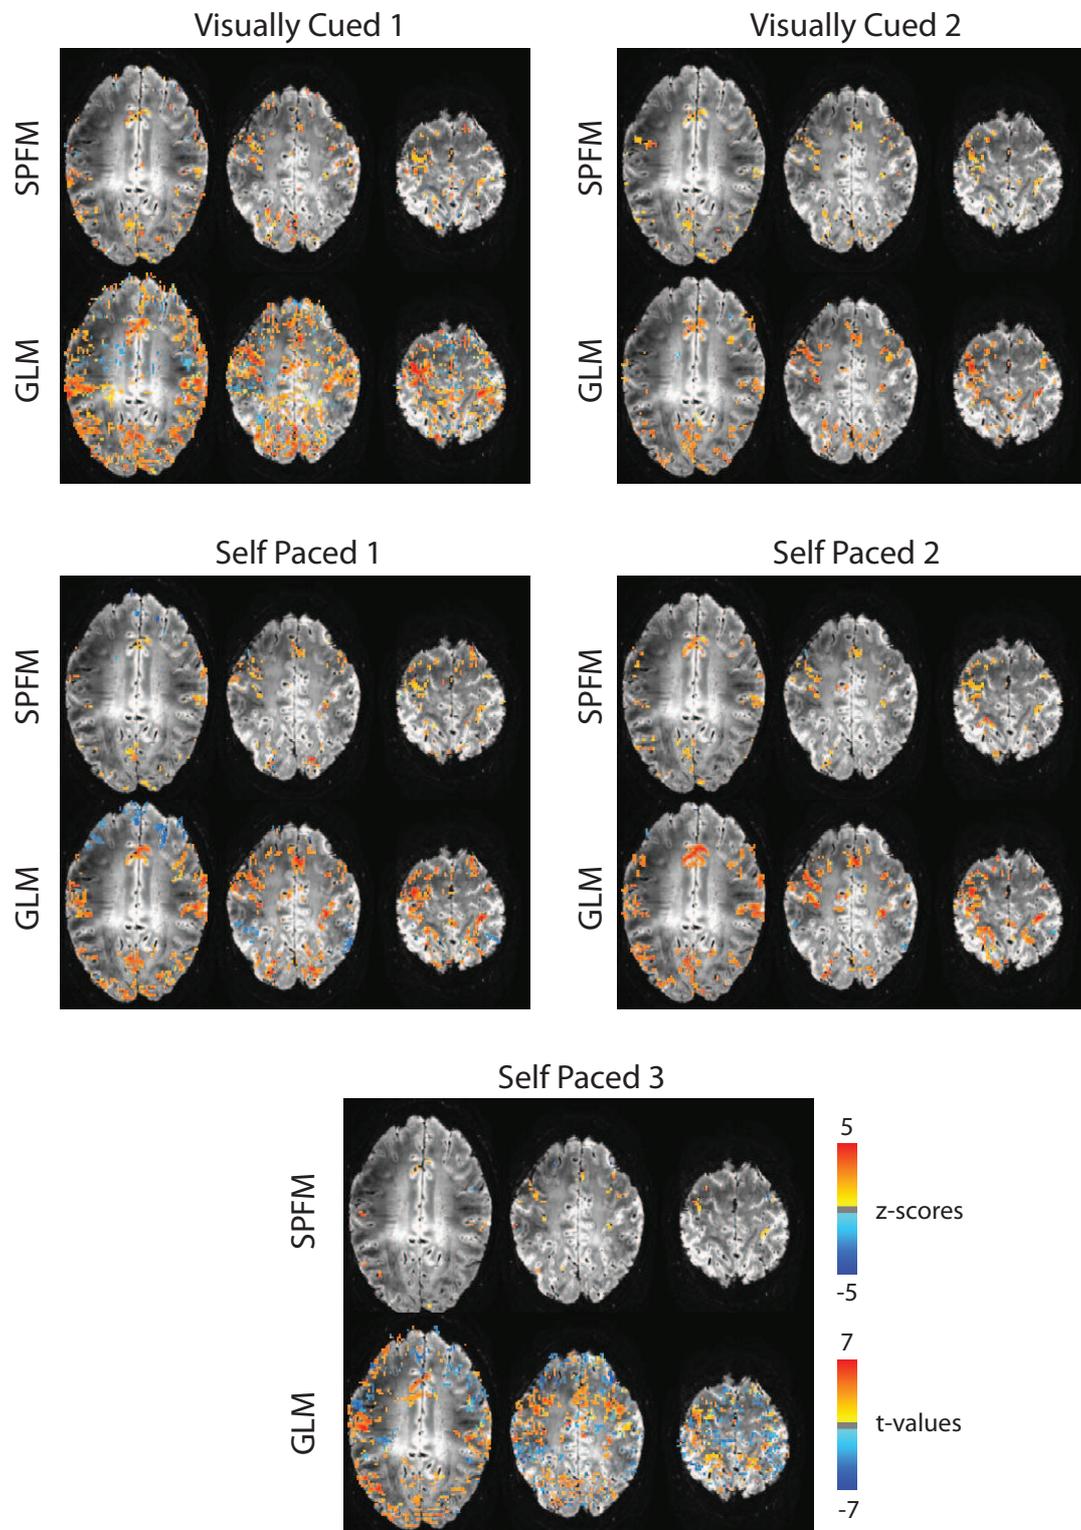


Figure 5.8: Comparison between SPFM activation maps and GLM maps for the two visually cued and three self paced finger tapping events of Subject A. SPFM maps display z -scores, whereas GLM maps display t -values.

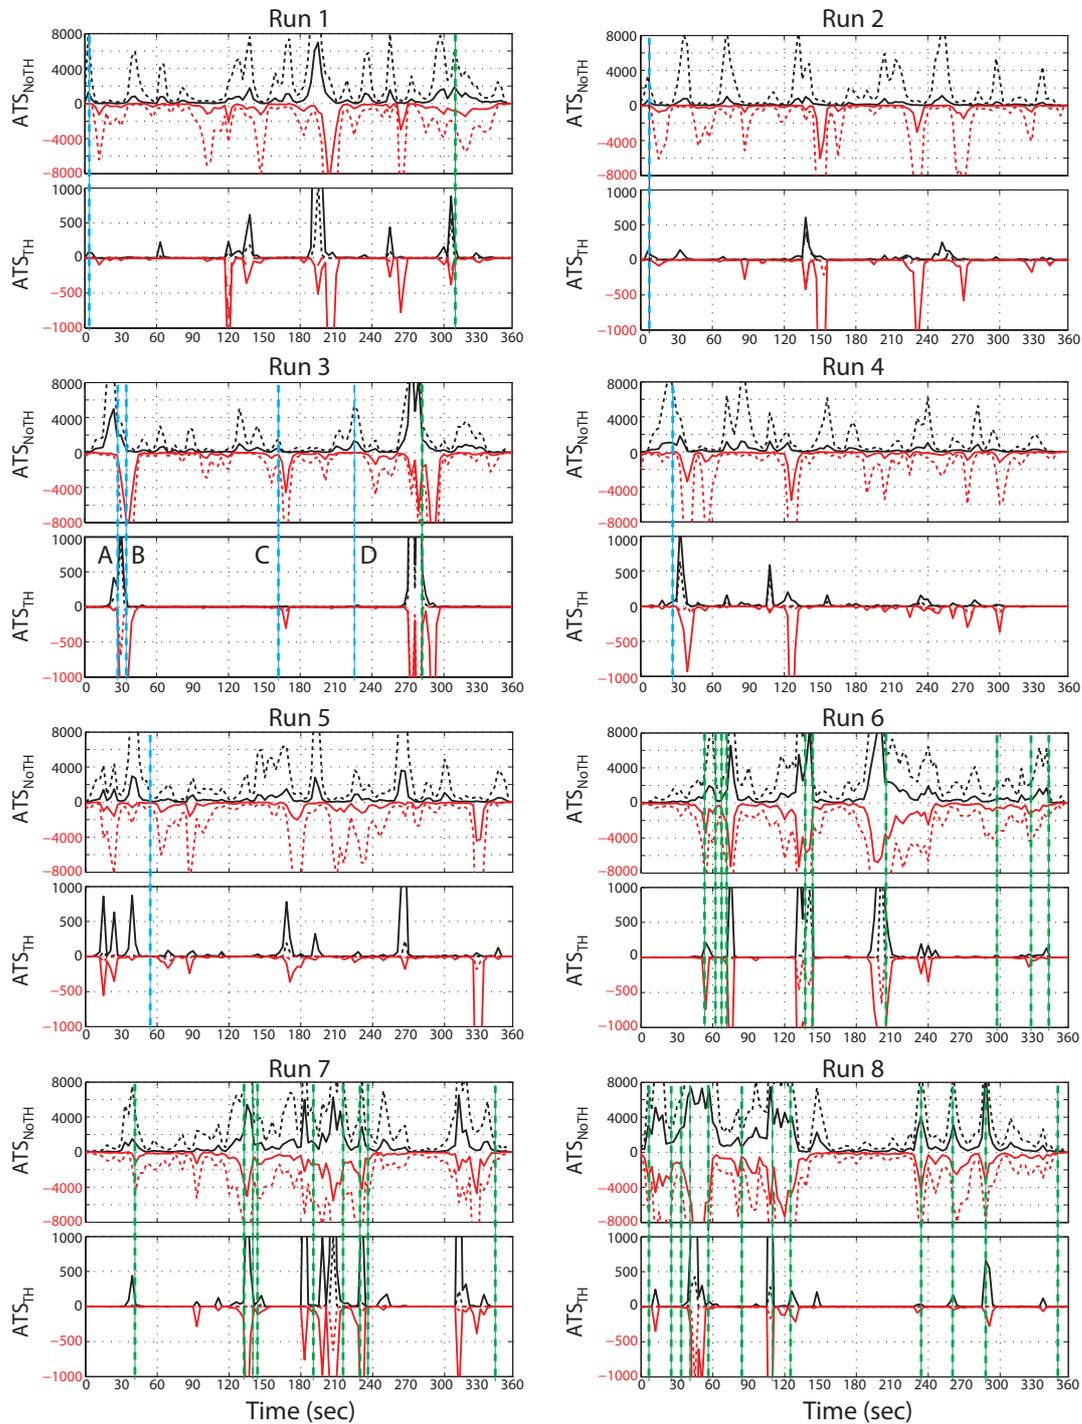


Figure 5.9: Activation Time Series computed with SPFM using the DS for the 8 epileptic datasets. Each plot shows the fMRI Activation Time Series before thresholding (top) and after thresholding ($Q < 0.05$, FDR-corrected) (bottom) computed with the MDL (solid line) and the UT criteria (dotted line). Blue dashed lines indicated times of IED and green dashed lines indicated times of significant head movement ($|d'| > 0.5$ mm/scan).

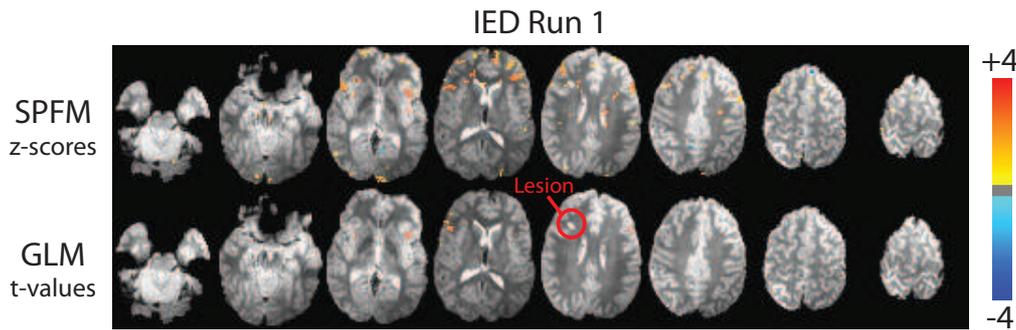


Figure 5.10: SPM and GLM maps of IED in Run 1.

supramarginal gyrus. The EEG-based GLM analysis of the IED ($Q < 0.05$, FDR-corrected) revealed fewer clusters of activations located in the inferior and middle frontal cortex, insular gyrus and caudate nucleus, either overlapping or adjacent to SPM clusters. No negative activations were observed in GLM maps (Figure 5.10).

SPFM allowed us to detect additional events in the fMRI signal. At 63s, positive BOLD signal changes were found bilaterally in areas of the frontal lobe with a more anterior spatial distribution of the activations in the same hemisphere as the lesion (shown in Figure 5.10). The events at 120 s and 138 s showed substantial positive and negative signal changes with patterns of activations which resemble sudden signal changes which might occur due to scanner or movement artefacts. The spatial distribution of the activations found in the event starting at 192 s showed a great resemblance with the maps of the EEG-marked IEDs A & B in Run 3, and the one in Run 4. This event showed positive signal changes in thalamus, caudate nucleus, frontal areas, entorhinal cortex and later deactivations in areas of the Default Mode Network (DMN) such as posterior cingulate, precuneus, superior and lateral parietal cortex, and frontal areas. The event found at 255 s showed negative signal changes in anterior orbitofrontal cortex bilaterally, and positive activations in entorhinal cortex, thalamus and occipital lobe. At 264 s, negative activations were seen at posterior cingulate, superior and lateral parietal cortex, occipital lobe and precuneus. Even though the event at time 306 s was simultaneous to large head movement with $|d'| = 2mm/scan$, which generated very distinctive spiking BOLD signal changes in the inferior orbitofrontal cortex bilaterally, later SPM maps showed significant ($Q < 0.05$, FDR-corrected) clusters of positive activation in frontal regions and thalamic areas of the same hemisphere as the lesion, whereas a negative cluster was also found symmetrically in thalamic areas.

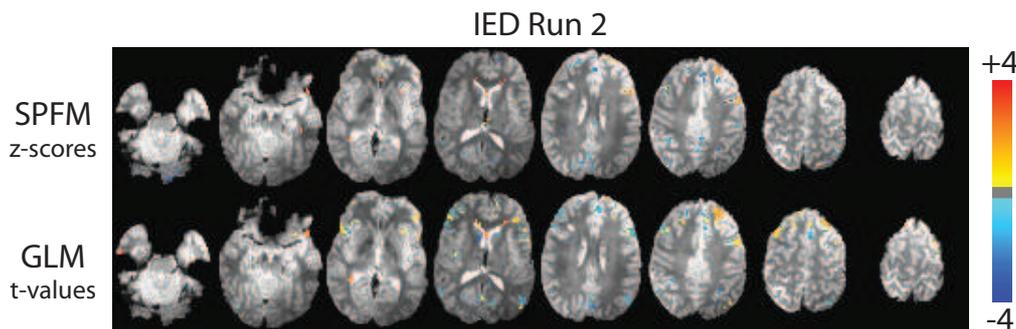


Figure 5.11: SPFM and GLM maps of IED in Run 2.

Run 2

One IED was identified at time 8.4 s. This IED was identified in both non-thresholded and thresholded ATS. The SPFM Activation Map of the IED indicated initial positive activations in frontal areas (inferior, middle and superior frontal gyri) and posterior lateral parietal cortex (supramarginal gyrus) and occipital lobe. Positive BOLD responses were later observed in the insular gyrus and anterior inferior temporal gyrus, tail of the caudate nucleus, fusiform gyrus and entorhinal cortex, concurrently with negative BOLD responses in the posterior cingulate cortex and lateral and superior parietal cortex. In subsequent maps, negative responses further extended into occipital lobe, internal capsule, caudate nucleus and the frontal lobe. In addition, positive responses were detected in the head of the caudate nucleus surrounding the anterior part of the lateral ventricles. The sensitivity of the EEG-based GLM was higher than SPFM and the analysis revealed significant clusters of positive and negative responses in the frontal lobe and the insular gyrus, positive responses in the inferior temporal gyrus and negative responses in the posterior cingulate gyrus, and middle and lateral posterior parietal (Figure 5.11).

Regarding the additional events detected by SPFM, positive activations were seen in frontal areas, caudate nucleus, precuneus and insula at 30 s. Negative activations were seen in the occipital lobe and precentral gyrus at 87 s. The event at 138 s showed positive activations in frontal and posterior areas in the same hemisphere and axial slices as the lesion, and negative activations in dorsal regions bilaterally and superior to the lesion. Later at 150 s, negative activations were predominantly found in DMN areas, bilateral insula but also positive activations in the entorhinal cortex bilaterally. Negative activations were observed in frontal areas, insula, posterior cingulate and parietal cortex at 231 s and 270 s, whereas at 330 s negative activations were located in large veins of the occipital-parietal and cerebellar-occipital sulci.

Run 3

Four IEDs were identified in this run at times 26.6 s (A), 32.0 s (B), 159.5 s (C) and 226.9 s (D). The first two IEDs (26.6 s and 32 s) were considered together since BOLD signal changes due to both IEDs overlapped in time. The EEG-based GLM statistics were adapted accordingly in order to consider both events (t -test of the two canonical HRF regressors and F -test of the 6 corresponding regressors). The SPM technique captured IEDs A and B and negative responses of IED C with thresholding, whereas IED D was only detected in the ATS without thresholding.

The sequence of BOLD signal changes associated with IED A and B was plotted in order to picture the dynamics of epileptiform network (Figure 5.14). The maps plot the signal fitted by the SPM method (canonical HRF model) from 21 s (-5 s before IED A) to 60 s (28s after IED B). Initial positive BOLD changes were predominantly seen in the dorsal frontal lobe, posterior superior frontal gyrus, posterior lateral parietal cortex and the insula. Interestingly, negative responses were also observed at this initial stage close to the lesion, middle part of the insula and superior to the main cluster of positive activation in the lateral parietal cortex. As illustrated in the figure, the timing of these BOLD changes preceded the time of the IED marked in the EEG. Later, positive responses were revealed in the thalamus, caudate nucleus, entorhinal cortex, fusiform gyrus, anterior parts of the inferior occipital lobe, and superior cerebellum. At 36s, negative responses began to appear in the posterior cingulate gyrus, precuneus and caudate nucleus. A predominant cluster of negative responses was found in the same areas of the lateral parietal cortex, which exhibited positive BOLD changes in earlier maps. Finally, negative responses extended into posterior parietal areas and the frontal lobe which peaked at 42 and 45 s. The compact SPM Activation Map of the activations corresponding to both IEDs is shown in Figure 5.12. The EEG-based GLM analysis of IEDs A and B revealed positive responses in the thalamus, caudate nucleus, dorsal frontal areas, entorhinal cortex and fusiform gyrus. Negative responses were predominantly seen in posterior brain regions with scattered frontal clusters. The insula showed both positive and negative responses in different clusters as detected with both methods.

The SPM maps for IED C found positive activations in dorsal inferior frontal areas and the insula. Negative responses were also observed in superior dorsal frontal areas, superior posterior frontal cortex close to central sulcus, superior parietal cortex, posterior cingulate, precuneus, and occipital lobe. In contrast, the GLM analysis only revealed negative responses in posterior areas of the occipital lobe with very

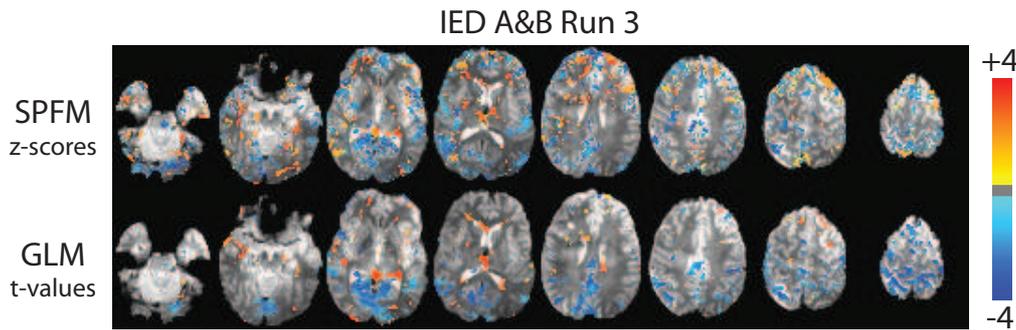


Figure 5.12: SPFM and GLM maps of IED A and B in Run 3.

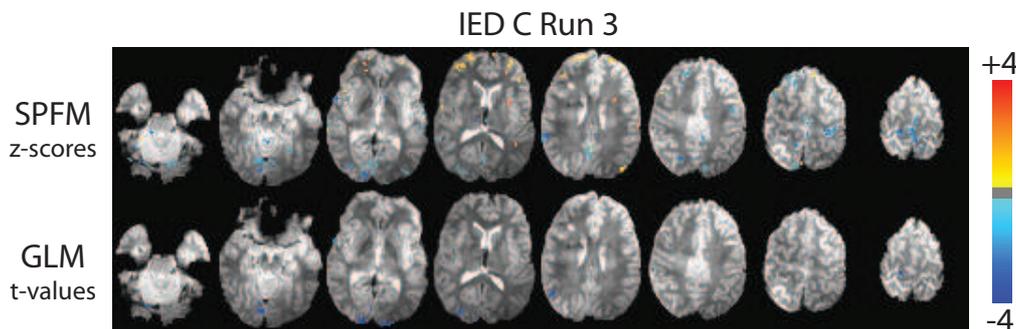


Figure 5.13: SPFM and GLM maps of IED C in Run 3.

little agreement with the SPFM results (Figure 5.13).

SPFM analysis of IED D found positive responses involving the caudate nucleus, insula, thalamus, caudate nucleus, dorsal superior frontal gyrus, middle temporal gyrus and, interestingly also positive responses in the lateral and superior parietal cortex and the occipital lobe were found in this IED. No voxels exceeded the established threshold in the GLM analysis (the minimum Q -value was 0.1221, FDR corrected). Therefore, for this IED we considered to reduce the GLM threshold until both SPFM and GLM maps presented equal number of significant voxels. Then, the GLM map revealed positive responses closely matching the areas found with SPFM (Figure 5.15).

The activation events detected by SPFM from 270 s to 294 s displayed a chaotic spatial distribution of positive and negative activations involving the entire cortex. Significant head movement $|d'| = 2\text{mm}/\text{scan}$ was detected at 282 s as highlighted in the corresponding ATS of Figure 5.9.

Run 4

One IED was identified at time 30.2 s, which was clearly seen in the thresholded ATS. Positive activations were initially observed in the insular cortex, anterior inferior frontal gyrus, anterior middle frontal gyrus, superior frontal gyrus, orbitofrontal

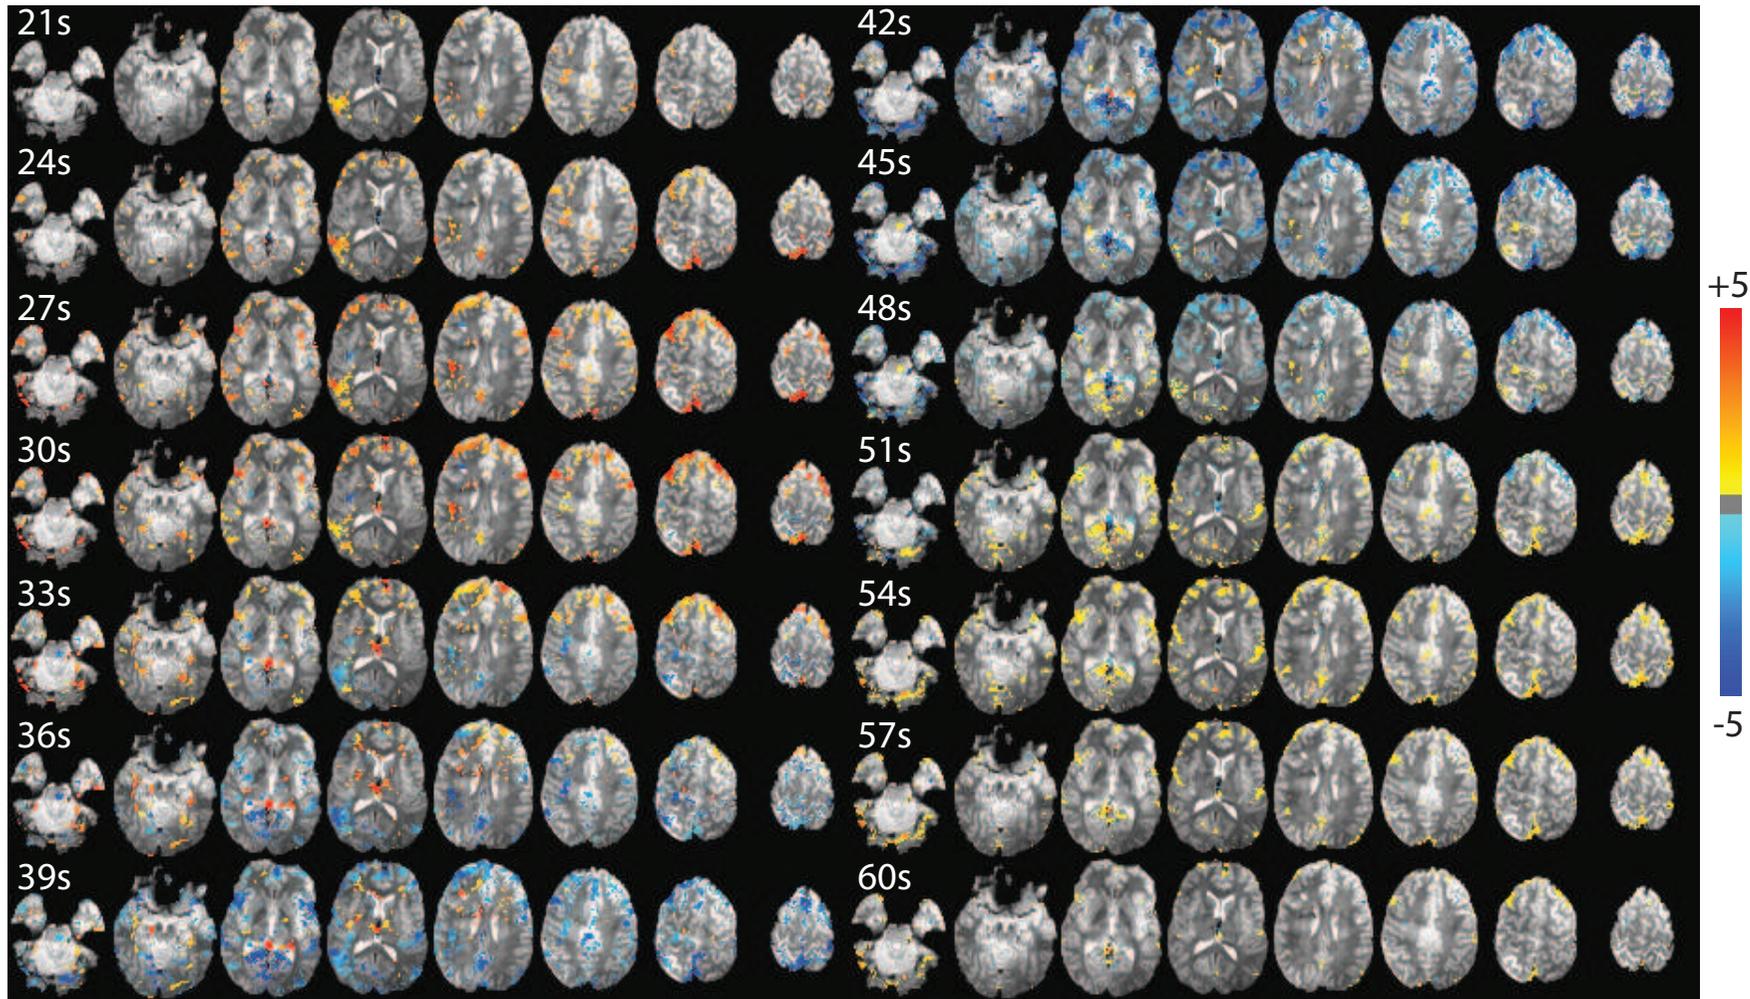


Figure 5.14: Sequence of SPM maps corresponding to the IED A and B in Run 3 (time of IED in EEG was 30.2 s). The sequence of images depicts the fitted signal (convolution of the canonical HRF with the signal estimated by the SPM method using the DS algorithm and the MDL criteria after debiasing) in scale of percentage signal change.

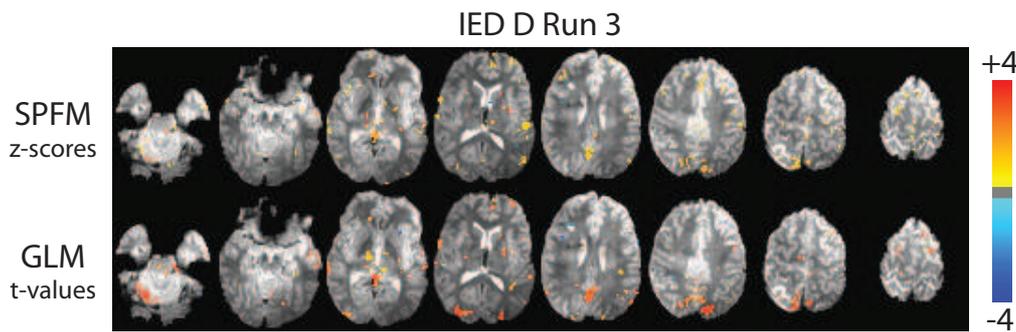


Figure 5.15: SPFM and GLM maps of IED D in Run 3.

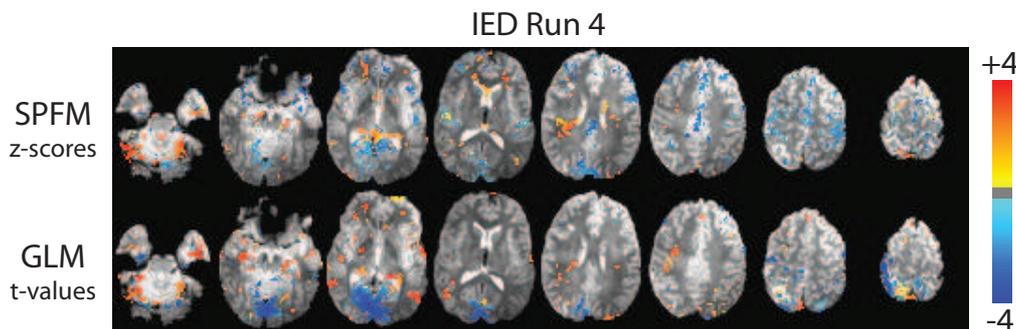


Figure 5.16: SPFM and GLM maps of IED in Run 4.

gyrus, middle cingulate gyrus, fusiform gyrus. Later positive activations were found in the head and tail of caudate nucleus, thalamus, posterior insula, entorhinal cortex, lateral parietal cortex, posterior superior temporal gyrus, superior parietal cortex, anterior inferior temporal gyrus, and the superior cerebellum. At the same time, the first negative responses appear in the precuneus and precentral gyrus, which later extended laterally into parietal areas, dorsally into the middle cingulate gyrus and the superior frontal lobes, and the occipital lobe. The areas revealed with EEG-based GLM were overlapping or adjacent to those ones found with SPFM except the large negative cluster found in the occipital lobe which was scattered in SPFM (Figure 5.16).

In this run SPFM also detected an event at time 108 s with positive activations in anterior frontal areas, anterior insular gyrus, superior parietal cortex along with negative activations in the posterior insular gyrus. At time 123 s, positive activations were observed in white matter voxels surrounding the head and tail of the lateral ventricles. This was followed at 126 s by negative activations in voxels near anterior and posterior sagittal sinus, cingulate cortex, insula and posterior areas of the superior frontal gyrus.

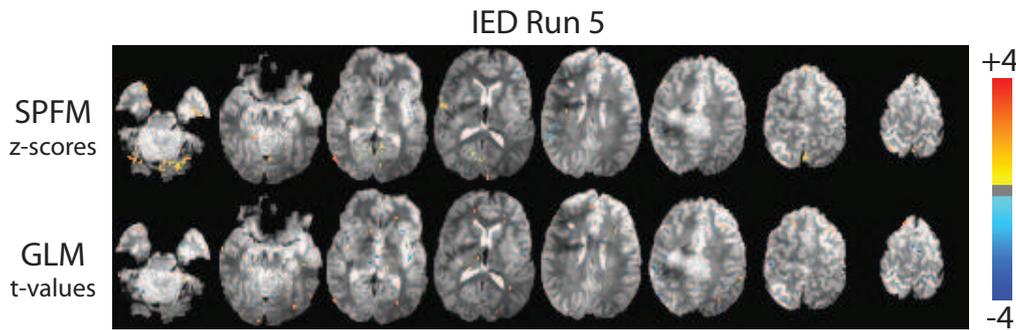


Figure 5.17: SPFM and GLM maps of IED in Run 5.

Run 5

One IED was identified at time 52.2 s. SPFM and EEG-based GLM analysis did not find any significant activation with $Q < 0.05$, FDR corrected for the EEG-marked IED in this run. Nevertheless, without thresholding the SPFM analysis showed clusters of positive responses in the entorhinal cortex, superior frontal gyrus, insula and inferior anterior temporal gyrus and in the main veins and sinus draining from the occipital lobe. Negative responses were observed in the dorsal medial frontal cortex, thalamus and precuneus. When reducing the threshold in GLM to match the number of voxels detected with SPFM ($p < 0.02$, uncorrected), the maps did not show overlapping areas except in entorhinal cortex and superior frontal gyrus, and negative responses in the thalamus, anterior inferior frontal gyrus and anterior part of the insular gyrus (Figure 5.17).

The fMRI data of this run exhibited large positive and negative fluctuations (signal changes larger than 10%) in WM areas of the same hemisphere as the lesion, corresponding to the peaks in the thresholded ATS at times 12 s, 21 s, 36 s, 66 s and 87 s. The events shown in the thresholded ATS at 189 s and 284 s consisted of positive activations in the posterior cingulate cortex, lateral parietal, precuneus and occipital lobe. The same regions were found to deactivate at time 327 s.

Runs 6, 7 and 8

No IEDs were detected in these runs in the EEG. These datasets exhibited considerable head motion (mean/std/max of $|d'|$ were 0.238/0.376/2.516 in run 6, 0.195/0.3/1.725 in run 7, 0.217/0.352/1.838 in run 8) with 35 head jerks in total, and SPFM detected significant responses in 30 of them. Across all runs, the SPFM maps at the time of head jerks displayed a distinctive pattern of positive and negative activations, easing their identification as motion-related events as shown in Figure 5.18.

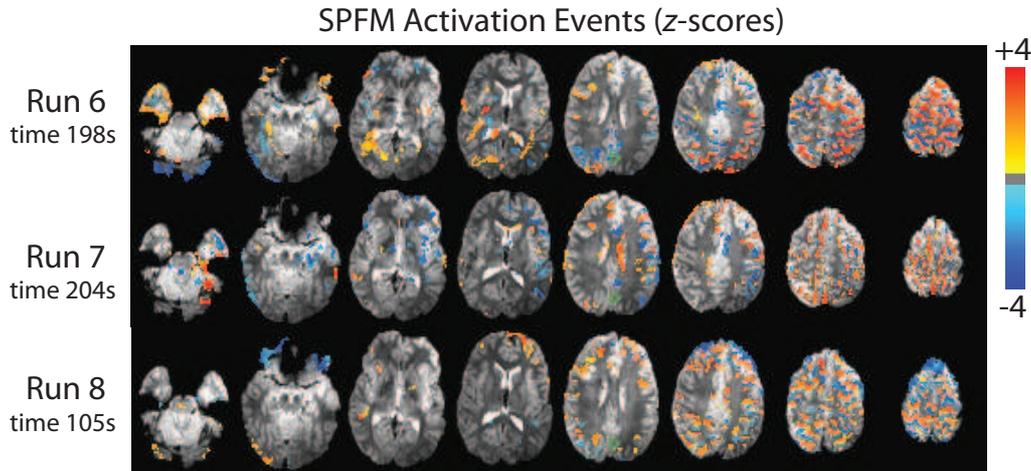


Figure 5.18: Three examples of SPFM Activation Events related to significant head motion $|d'| > 0.05$ mm/scan in runs 6, 7 and 8.

5.4 Discussion

Our results demonstrate that sparse estimation techniques are an adequate tool to detect single trial events in a paradigm-free method where the timing of the activations is unknown. We made use of L_1 -norm regularized estimators to deconvolve the signal in time and reveal significant responses according to an HRF shape. This contrasts with the manner in which sparse estimators are typically used for fMRI data analysis in the spatial dimension of spatio-temporal models [354, 121, 224] and/or feature selection [307], even though some attempts have been presented to consider temporal sparse events in fMRI [191] and EEG/MEG [45]. It has recently been demonstrated that spatial sparsity is the quality which drives spatial ICA decompositions rather than independence [89, 208].

Methodological issues and simulations

Consistent with previous theoretical results, sparser estimates were obtained with DS than with LASSO [180]. On average, however, there were no significant differences between both estimators in terms of the mean square error (MSE_s) as shown in Figures 5.5 and 5.6 when the regularization parameter is chosen based on the MDL, UT and LUT criteria which tend to compute sparser estimates than AIC, in agreement with the results in [36]. This effect was also observed in the specificity and sensitivity analysis shown in Figure 5.4. Although we found the use of L_1 -norm convex estimators very useful in our experiments, it has been proved that non-convex methods based on L_1 and L_2 -norm reweighted methods are superior

[366]. Using sparsity as a measure may be appropriate in the case of single-trial event-related responses, but arguable since the actual degree of sparsity of the signal is unknown in real situations. Additional simulations (data not shown) have demonstrated that the SPFM method using MDL and UT approaches to select λ can recover BOLD signal changes due to longer events. Furthermore, the LASSO and the DS estimators have been shown to achieve accurate estimates of relatively non-sparse signals [68, 180].

In this study, the regularization parameter λ was selected according to model selection criteria (MDL and AIC), following the work by Zou et. al. (2007) [384], and thresholding approaches (UT and LUT) [10, 101]. Our simulations demonstrated that none of the criteria investigated was superior in all scenarios and the selection depends on tSNR and sparsity of the signal. Our simulation results illustrated the well-known tendency of AIC to overfit the data when trying to minimize the prediction error, and a superior performance of MDL which tends to select the true model when this is included among the candidate models [384]. In practice, the use of the MDL and UT selection methods are recommended since they offered better false positive control. The experimental results demonstrated the superiority of the MDL criteria in the detection of sparse BOLD events. Furthermore, homogeneous maps of λ were obtained for all the selection criteria in both experimental paradigms even though voxelwise selection of λ was used in this study. This may potentially cause overfitting of the data and future studies will also consider alternative approaches where the same value of λ is used for a neighbourhood of voxels, e.g. according to tissue classification.

Uncorrelated Gaussian noise was assumed in this thesis. However, SPFM could deal with correlated noise via data prewhitening providing the serial correlations were accurately estimated [373]. Yet, our simulations tested the performance of the method with additional sinusoidal noise terms, representing cardiac and respiratory fluctuations. The results demonstrated better performance of the estimators in that scenario, similar to observations in [180], suggesting that sparse estimation can be seen as an accurate band-pass filtering approach tuned to the HRF activations which is able to denoise the BOLD signal [204] as was illustrated with the simulations in Figure 5.3. In that sense, the DS and LASSO estimators showed robustness against HRF variability [4, 103, 162, 242, 323] in the estimate of the simulated BOLD responses (Figures 5.5 and 5.6) in contrast to the traditional degradation of the LS estimator used in GLM with model mismatches especially at high SNR [64, 211].

We also evaluated the specificity and sensitivity of the technique to detect event-related responses in time. In fMRI data analysis, false positive and false negative ratios have typically been measured over space assuming a model of experimental effects. However, a temporal model of the events is not available in PFM analysis and so sensitivity and specificity need also be measured in time. In this study, both measures were computed regardless of the standard error of the estimates in the simulations. This was done because sparse estimation techniques may potentially reveal BOLD activations based on the coefficient estimates [191, 379]. Even so we also aimed to provide statistical inference of the estimates and this was done based on the debiasing model and conditional on the solution found by the DS and LASSO estimator [120]. Mathematically speaking, since the individual coefficient estimates will typically exhibit a probability distribution with more condensation of zeros estimates, they may be far from normally distributed. This suggests that summarizing uncertainty by standard errors and standard t - and F -tests may not be completely appropriate and this topic deserves further investigation [272, 339]. Therefore, alternative methods based on the residual diagnostic statistics may also become useful for statistical model validation and assessment [229, 232, 297].

Experimental data

Visuomotor paradigm

The analysis of the visuomotor paradigm with SPFM confirmed the activations detected with the PFM method described in Chapter 4. In Chapter 3, paradigm free mapping was initially proposed based on the deconvolution of the BOLD signal using Ridge Regression, and statistical inference required the definition of a baseline period in order to compute the t -statistics. Here, this requirement is not necessary by virtue of the subset selection features of L_1 -norm regularized estimators. The fact that no posterior thresholding is required when using sparse estimation algorithms to deconvolve the signal suggests that this method eases the detection of single trial fMRI BOLD responses and facilitates the interpretation of results without decreasing temporal specificity and sensitivity to detect the finger tapping events (compare Figures 5.7 and 4.9).

Spatial averaging of the deconvolved data was performed in PFM in order to compute the t -statistics, considering a neighbourhood of $L=5$ voxels. In contrast, the visuomotor datasets were analyzed with SPFM without spatial smoothing and this may explain why fewer voxels are detected with SPFM than with PFM. The

SPFM maps associated with finger tapping events showed functional activations in the same cortical areas as were detected with PFM (Figure 5.8) and we observed that 75 % of the voxels found active in SPFM were also active in the FDR-corrected GLM maps.

Epilepsy data

We also evaluated the potential application of SPFM to detect IED in 8 fMRI runs acquired in a patient with idiopathic generalized epilepsy. SPFM was able to detect 7 of the 8 IEDs marked in the EEG data without the need to specify the timing of the IED. In general, the results of the analysis showed good visual concordance between the SPFM maps and the EEG-based GLM maps, particularly in those IEDs that were detected after FDR correction.

The sequential analysis provided by SPFM maps suggest that positive BOLD signal changes precede negative ones (Figure 5.14) assuming that the HRF model is constant across regions. Positive activations were mainly located in the thalamus, the insular gyrus, frontal areas, limbic structures in the medial superior temporal cortex, superior cerebellum, and borders of lateral ventricles. Involvement of thalamic structures in the generation of generalized epileptic discharges has been consistently found in EEG-fMRI studies [2, 147, 254]. Positive and negative activation in the insula might be explained by its connections to the thalamus [2, 11, 161], and similarly for cerebellar activations [147]. With SPFM we also observe positive activations in frontal areas which achieved maximum amplitude at the initial stages of the IED BOLD changes. Bursts of generalized waves were also observed in the EEG in frontal and central regions, but are not generally seen in fMRI data [2]. Activation in limbic structures such as the entorhinal cortex have been found with thalamic stimulation in animal studies of epilepsy [3]. Positive and negative activations in areas of the occipital lobe in generalized epilepsy have been found in [2, 11, 161]. Positive responses in the borders of the lateral ventricles were unexpected. These activations are hard to explain due to BOLD effects and are generally associated with physiological fluctuations or head motion. No physiological recordings were measured to test the first explanation, and no significant head motions ($\|d'\| > 0.5$ mm/scan) were observed at the times of the epileptic discharges. However, activations in WM areas are believed to result from sudden changes of ventricular volume due to the general increase in blood volume, thereby changing the NMR signal from these voxels [147].

Negative responses were predominantly seen in brain areas of the default mode network (DMN) [53], such as the precuneus, posterior cingulate, bilateral posterior parietal cortex, anterior cingulate cortex and anterior frontal regions at a later stage. The deactivation of the DMN has been revealed in several EEG-fMRI studies of generalized epilepsy [2, 147, 161, 256, 308], suggesting a temporal suspension of normal resting state brain function [196]. The amplitude of negative signal changes was of the same magnitude as positive activations [196]. Additional negative responses were sometimes revealed in the head of the caudate nucleus, posterior to the lateral horns of the lateral ventricles, as it has previously been described in some patients with idiopathic generalized epilepsy [2, 161, 254]. When negative responses were clearly observed, these were always bilateral and never found solely in anterior regions [196].

Apart from the activations related to the IEDs, SPFM allowed us to identify additional activation events which were clearly marked in the ATS but not in the EEG, and that were not concordant to significant head motion. The spatial distribution of the activations of some of these events showed strong resemblance with the IED-related activation maps, e.g. event at 192 s in Run 1, or involved cortical areas commonly seen in idiopathic generalized epilepsy, such as the caudate nucleus or the thalamus. Cautiously, the maps of other events shown in the SPFM ATS showed significant and synchronous positive or negative BOLD signal changes across extensive areas of the cortex, particularly near cortical sulci and sinus, which might be related to cardiac pulsations. The long TR used to acquire these epileptic datasets (TR of 3 s) causes cardiac and respiratory fluctuations of the signal to overlap in frequency with the BOLD responses. Therefore, we cannot neglect the possibility that some of the activations detected with SPFM might be associated with physiologically-related signal changes since taking good account for cardiac and respiration-related fluctuations has been proven decisive in resting state fMRI studies [72] and also in epilepsy [160, 216, 351]. Ultimately, it will be necessary to reexamine the EEG data at these time points in order to elucidate whether these events are related to some type of epileptiform activity.

In this study, SPFM allowed us to map sequential BOLD signal changes related to interictal discharges (Figure 5.14), but the method could also be applied in ictal studies [12, 179, 197, 348] if the patient experienced partial seizures during clinical assessment with fMRI. Sequential analysis of ictal fMRI data has successfully allowed the identification of the brain structures involved in seizure generation and map the

dynamics of the epileptiform network during pre-ictal, ictal and post-ictal periods with temporal resolution of 10 s [16, 99, 100]. As a step forward, SPFM allows the visualization of the sequence of activations with the actual temporal resolution of the fMRI acquisition. A paradigm free mapping approach integrating EEG and fMRI models has been recently proposed based on the Kalman filter [97], but this method find difficulties in tracking sharp spatio-temporal sparse events such as the ones observed in interictal seizures. Similarly, the modulus operandi of SPFM can be combined with continuous EEG source imaging techniques [357, 358, 359].

The canonical HRF was used as a haemodynamic response model, but variability in the BOLD response to IED may exist between areas of the epileptic network and subjects [60]. FIR or Fourier basis sets can deconvolve the HRF shape with prior information of the stimulus onset [60] or after obtaining this information with ICA [210]. To overcome this problem, Bagshaw et al. (2004) proposed mapping the maximum of the t -values of 4 different regressors created with the convolution of four HRF shapes with peaks at 3s, 5s, 7s and 9s [17]. According to our simulations, SPFM using the canonical HRF is robust against this range of mismatches in terms of MSE of the haemodynamic component of the signal. Besides, the canonical HRF has accurately modelled activations (deactivations) in the most relevant areas involved in idiopathic generalized epilepsy, such as the thalamus, caudate nucleus and default mode areas [254].

Head motion is a major problem in clinical fMRI studies and we found measuring the absolute scan-to-scan displacement from the realignment parameters very useful to detect significant head jerks [209]. Similar to temporal clustering analysis [160], our results demonstrated that SPFM is very sensitive to movement-related effects since activation events were found in 85 % (30/35) of the head jerks marked with our movement-threshold ($|d'| > 0.5$ mm/scan) in the epilepsy study. Therefore, adding the 6 realignment parameters as nuisance covariates into the design matrix prior to debiasing, as it is routinely done in standard fMRI analysis [137], seems to be insufficient to model motion-related signal changes of such a high magnitude. Alternative procedures can be considered to deal with this problem. To begin with, 24 instead of 6 motion-related regressors can be added into the design matrix based on a Volterra expansion of the realignment parameters [137, 209, 304, 357, 359]. In addition, one might also consider including scan-nulling regressors marking scans with significant head jerks [209]. Note, however, this might not be practical for datasets with reduced numbers of scans since it will considerably reduce the de-

degrees of freedom of the statistics. Motion regressors could also be incorporated into the model prior to the DS or the LASSO estimation, which would then select the most representative subset of HRF or motion-related regressors. Nevertheless, this formulation causes considerable reduction in sensitivity as additional analyses have demonstrated (results not shown), calling for the design of mixed methods with L_1 - and L_2 -norm terms [198].

Taking those points into consideration, these initial results applying SPFM to study interictal epileptic discharges are encouraging in that SPFM can provide complementary information about when to perform EEG source localization, and not only that the EEG informs about how to define a GLM analysis as it has been the common practice in EEG-fMRI epileptic studies [148, 358]. Modelling only IED detected on the scalp EEG may not reflect the abundant underlying epileptic activity recorded with intracranial recordings [61, 358]. Finally, SPFM may help to retrieve datasets where no IED are evident in the EEG data, similar to other model free approaches such as ICA [240, 254, 304] or temporal clustering analysis (TCA) [160, 258, 256]. All in all, SPFM can help to enhance our understanding of epileptiform activity and address some of the questions of inter-subject and intra-subject variability observed in EEG-fMRI studies of epilepsy [148, 357].

5.5 Conclusion

This chapter described a PFM method based on sparse estimation techniques, Sparse Paradigm Free Mapping. In particular, the LASSO and DS estimators were investigated in this study, in combination with model selection and thresholding approaches to select the regularization parameter. The results of our simulations demonstrated that the method is able to provide higher temporal specificity and sensitivity in the detection of event-related BOLD responses with an accurate HRF model. Even so, the benefits from using sparse estimators are better seen in terms of its denoising capability, where the technique is able to extract BOLD signal changes at low SNR conditions and with HRF mismatch. Notably, the method showed preliminary, but promising, results in the study of epileptic activity with fMRI without relying on EEG information.

The next chapter will also rely on sparse estimation techniques in order to describe a multicomponent PFM approach, which decomposes the fMRI signal into multiple sources.

Chapter 6

Multicomponent Paradigm Free Mapping

This chapter extends the principles of paradigm free mapping into a semiparametric formulation which allows the decomposition of the fMRI signal into multiple components under a generalized linear model: Multicomponent Paradigm Free Mapping. The technique relies on the algorithm of morphological component analysis and enables both the paradigm-free detection of event-related single-trial BOLD responses and automatic extraction of baseline fluctuations in the fMRI signal (§6.1). The operation of the technique is evaluated with simulations and real fMRI data. Its usefulness is illustrated by comparing it with RETROICOR, since the baseline component also comprises physiological cardiac and respiratory fluctuations which are extracted without the need of external physiological recordings (§6.2 and §6.3). Finally, these results and possible extensions of the technique are discussed in §6.4.

6.1 Theory

6.1.1 Signal Model

Previous PFM approaches have assumed the following signal model of the fMRI voxel time series:

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \boldsymbol{\epsilon}, \quad (6.1)$$

where \mathbf{y} is a N -length column vector representing the fMRI voxel time series, $\mathbf{H}\mathbf{s}$ comprises the haemodynamic component of the signal resulting from BOLD signal changes as a response to stimuli or neural events, and the noise term $\boldsymbol{\epsilon}$ includes

baseline fluctuations and noise components of the signal. Among other sources, baseline fluctuations in fMRI include scanner drifts [336], cardiac and respiratory physiological fluctuations [38, 40, 71, 86, 90, 144, 320, 365] or movement artefacts [137, 209]. Importantly, baseline fluctuations of systematic origin have been observed not only in-vivo [35], but also in cadavers [321] and phantoms [228].

In fMRI data analysis, one faces a choice when dealing with noise fluctuations. On the one hand, they can be removed prior to the statistical analysis by means of detrending methods. On the other hand, descriptive time courses of the fluctuations can be incorporated as additional covariates of the linear model and estimated simultaneously with the stimulus effects [188, 229]. It is important to realize that these approaches are not equivalent and by using a detrending approach one assumes that any potential correlation between the fluctuations and the effects is removed and neglected in the statistical analysis [219, 251].

Similar to [116, 231, 251], we propose extending (6.1) to a partially linear model

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{g} + \boldsymbol{\epsilon}, \quad (6.2)$$

where \mathbf{g} is a column vector of length N which describes a nonparametric and arbitrary function, modelling the systematic and physiological fluctuations of the noise, and $\boldsymbol{\epsilon}$ is random white noise with variance σ^2 . Furthermore, we assume that the nonparametric component can be described by a linear combination of basis functions within a complete dictionary, i.e. $\mathbf{g} = \boldsymbol{\Phi}\boldsymbol{\alpha}$, such that (6.2) can be written as

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \boldsymbol{\Phi}\boldsymbol{\alpha} + \boldsymbol{\epsilon}. \quad (6.3)$$

Following [115] a dictionary $\boldsymbol{\Phi} = [\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_L]$ is defined as a $N \times L$ matrix whose columns are basis functions with unit euclidean norm, i.e. $\|\boldsymbol{\phi}_l\|_2^2 = 1$, for $l = 1, \dots, L$. When $L = N$, we are dealing with a complete dictionary, while overcompleteness occurs when \mathbf{g} is represented by more coefficients than the number of samples, i.e. $L > N$.

The main contribution of Multicomponent PFM is that both the neural-related haemodynamic effects and the baseline fluctuations are estimated from the data in a paradigm-free scenario. This contrasts with existing approaches in fMRI data analysis which, based on a semiparametric model, propose decomposing the fMRI voxel time series into a parametric part assuming prior knowledge of the stimuli or experimental effects, and a nonparametric baseline term [116, 231, 251].

6.1.2 Morphological Component Analysis

Assuming the model in (6.2), we face solving an underdetermined problem where the total number of unknown model coefficients $N + L$ is larger than the number of observations N . Consequently, similar to the LASSO and DS estimators investigated in Chapter 5, we need to rely on sparsity priors on \mathbf{s} and $\boldsymbol{\alpha}$ so that the decomposition of the signal becomes feasible. Here, the morphological component analysis (MCA) algorithm [330, 331] is used in order to decompose the fMRI signal into two components which are morphologically distinct, the haemodynamic component comprising the BOLD responses of the fMRI time series and baseline fluctuations. The term morphological component arises from the signal/image processing literature and means that the component has a distinctive morphological signature, e.g. piecewise smooth and spike components in a signal or different textures in an image, which facilitates its identification among other components of the signal.

The MCA algorithm solves the following constrained optimization problem

$$\{\hat{\mathbf{s}}, \hat{\boldsymbol{\alpha}}\} = \min_{\mathbf{s}, \boldsymbol{\alpha}} \|\mathbf{s}\|_1 + \|\boldsymbol{\alpha}\|_1 \quad \text{subject to} \quad \|\mathbf{y} - \mathbf{H}\mathbf{s} - \boldsymbol{\Phi}\boldsymbol{\alpha}\|_2^2 \leq \delta, \quad (6.4)$$

which is equivalent to minimizing the Lagrangian

$$\{\hat{\mathbf{s}}, \hat{\boldsymbol{\alpha}}\} = \min_{\mathbf{s}, \boldsymbol{\alpha}} \|\mathbf{y} - \mathbf{H}\mathbf{s} - \boldsymbol{\Phi}\boldsymbol{\alpha}\|_2^2 + \lambda\|\mathbf{s}\|_1 + \lambda\|\boldsymbol{\alpha}\|_1, \quad (6.5)$$

where the nonnegative, real regularization parameter λ balances the sparsity of the representation of the components in the dictionaries against the model imperfection measured in terms of residual sum of squares.

Consequently, the way in which MCA operates can be understood as a generalization of a GLM analysis where, based on an overcomplete representation of the fMRI signal and relying on the subset selection features of L_1 -norm estimators, the basis functions describing each component are estimated from the data instead of being specified in advance.

Two conditions are important to understand how the MCA algorithm operates [115]. First, the assumption that each component must be sparsely represented by its dictionary, i.e. \mathbf{s} and $\boldsymbol{\alpha}$ are sparse coefficient vectors where few coefficients are large and the rest can be neglected. This condition has been successfully investigated in Chapter 5 where the haemodynamic component of the signal due to event-related BOLD responses allowed a sparse representation in the HRF-based

dictionary \mathbf{H} . Second, when both dictionaries are considered together, each of them must efficiently represent a unique morphological component of the signal, and be highly inefficient in representing other components [115]. In other words, the representation of baseline fluctuations in the haemodynamic dictionary \mathbf{H} cannot be as sparse as it is in Φ . This condition is similar to the collinearity condition between covariates in linear regression and standard GLM analysis [9], and it can be mathematically understood as minimizing the correlation between the dictionaries, also known as mutual incoherence [115].

To solve (6.4) or (6.5), the MCA algorithm proceeds in a recursive fashion. Let $x^{(n)}$ denote the value of an arbitrary variable x at iteration n . If the component $\mathbf{H}\mathbf{s}^{(n)}$ is fixed, the estimation of the vector $\boldsymbol{\alpha}^{(n+1)}$ can be computed from the marginal residuals at each iteration, $\mathbf{r}_{\mathbf{s}^{(n)}} = \mathbf{y} - \mathbf{H}\mathbf{s}^{(n)}$, as follows

$$\hat{\boldsymbol{\alpha}}^{(n+1)} = \min_{\boldsymbol{\alpha}} \|\mathbf{r}_{\mathbf{s}^{(n)}} - \Phi\boldsymbol{\alpha}\|_2^2 + \lambda^{(n)}\|\boldsymbol{\alpha}\|_1. \quad (6.6)$$

Since most of the haemodynamic component has been removed, the LASSO estimator in (6.6) will select the most salient features to describe the baseline fluctuations. Equivalently, the sparse representation of the haemodynamic component $\mathbf{s}^{(n+1)}$ can be estimated as

$$\hat{\mathbf{s}}^{(n+1)} = \min_{\mathbf{s}} \|\mathbf{r}_{\boldsymbol{\alpha}^{(n)}} - \mathbf{H}\mathbf{s}\|_2^2 + \lambda^{(n)}\|\mathbf{s}\|_1, \quad (6.7)$$

where $\mathbf{r}_{\boldsymbol{\alpha}^{(n)}} = \mathbf{y} - \Phi\boldsymbol{\alpha}^{(n)}$.

Similar to the homotopy algorithms studied in Chapter 5, the algorithm starts with an extreme value for λ where the null solution with all components equal to zero is feasible, and hence the initial estimate of the residuals is equal to the fMRI time series, \mathbf{y} . The initialization of λ is given by the maximum correlation of the time series with both dictionaries, i.e. $\lambda_{max} = \max\{\|\mathbf{H}^T\mathbf{y}\|_\infty, \|\Phi^T\mathbf{y}\|_\infty\}$. Based on a block-coordinate relaxation algorithm [310], the value of λ is then reduced at each iteration according to a convergence strategy and the MCA cycles through each component so that the estimates of \mathbf{s} and $\boldsymbol{\alpha}$ are updated with the most salient features which describe each morphological component. The algorithm continues with decreasing values of λ until it is below a given stop criteria, λ_{min} . Note that if the algorithm continues until $\lambda = 0$, the entire regularization paths for both morphological dictionaries are computed.

An important computational consideration is that if the dictionary Φ is orthonor-

Algorithm 2: : Multicomponent Paradigm Free Mapping

Parameters: Convolution (Toeplitz) matrix \mathbf{H} according to the HRF model, dictionary for the nonparametric component Φ , stopping factor ν or λ_{min} , update strategy for λ .

For each fMRI voxel time series \mathbf{y} :

Estimate $\hat{\sigma}_{MAD}$ and set $\lambda_{min} = \nu \hat{\sigma}_{MAD}$

Initialization: $\hat{\mathbf{s}} = \mathbf{0}$, $\boldsymbol{\alpha} = \mathbf{0}$, $\lambda_{max} = \max(\|\mathbf{H}^T \mathbf{y}\|_\infty, \|\Phi^T \mathbf{y}\|_\infty)$;

while $\lambda(n) < \lambda_{min}$ **do**

 Compute marginal residuals for nonparametric part, $\mathbf{r}_{\mathbf{s}(n)}$;

 Update the nonparametric estimate by solving the LASSO in (6.6) (possibly via soft-thresholding);

 Compute marginal residuals for haemodynamic component $\mathbf{r}_{\boldsymbol{\alpha}(n)}$;

 Update the haemodynamic estimate by solving the LASSO in (6.7) with an L_1 -norm solver (e.g. a homotopy algorithm);

 Update the regularization parameter $\lambda^{(n)}$ according to a given strategy;

end

 Debiasing to reduce estimation bias of the morphological components.

mal, the LASSO coefficients α_l , $l = 1, \dots, L$, can be efficiently computed from the least squares (LS) estimates by soft-thresholding [233]

$$\hat{\alpha}_l = \text{sign}(\hat{\alpha}_l^{LS}) (|\hat{\alpha}_l^{LS}| - \lambda)_+ . \quad (6.8)$$

where $(a)_+ = \max\{a, 0\}$. Nevertheless, notice that the haemodynamic dictionary used in PFM approaches, \mathbf{H} , is not orthonormal. Therefore, computing the estimate of \mathbf{s} needs to be done with an L_1 -norm regularization solver. The same homotopy algorithm investigated in Chapter 5 is used in this study [15]. A pseudocode implementation of the MCA procedure in Multicomponent PFM is given in Algorithm 2.

6.1.3 Choice of the dictionaries

Choosing an appropriate dictionary to describe the baseline fluctuations is decisive in the successful operation of the MCA algorithm. A wide range of candidate dictionaries is now available to obtain sparse representations of a signal, such as wavelet packets, Gabor dictionaries, ridgelets, chirplets or cosine and sine packets [233, 306, 331]. The selection of the dictionaries is largely influenced by the type of data to decompose and prior knowledge of the possible components of the baseline fluctuations might help the choice of the dictionary to describe them. Importantly, one needs to minimize the mutual coherence between the dictionaries in the sense that the haemodynamic component of the signal cannot be more sparsely represented

in Φ than it is in \mathbf{H} . The main assumption in PFM is that the BOLD responses are transient, stochastic events for which no timing model can be assumed in advance. In that scenario of stochastic event-related responses, the energy of the haemodynamic component spreads across a wide range of frequencies [138, 217, 218, 219]. Therefore, the haemodynamic component does not admit a sparse representation in a Fourier dictionary, which then becomes a suitable choice to describe the baselines. Furthermore, an important source of the baseline variability of the fMRI arises from low frequency systematic drifts and physiological cardiac and respiratory rates [35], especially at high MR fields [343], which can be represented as single or linear combination of sinusoidal signals [144, 229].

Taking these points into account, we here suggest to use a dictionary Φ of size $N \times 2N$, i.e. $L = 2N$, with basis functions from the discrete cosine and sine transform (DCST)¹ such that $\boldsymbol{\alpha} = [a_0, b_0, \dots, a_{N-1}, b_{N-1}]^T$ and the nonparametric term of the signal, $\mathbf{g} = [g_0, \dots, g_{N-1}]^T = \Phi \boldsymbol{\alpha}$, is modelled as [233]:

$$g_n = \sum_{k=1}^N c_k \left[a_k \cos \left(\frac{k\pi}{N} \left(n + \frac{1}{2} \right) \right) + b_k \sin \left(\frac{k\pi}{N} \left(n + \frac{1}{2} \right) \right) \right], \quad n = 0, \dots, N-1, \quad (6.9)$$

where the coefficients of the discrete cosine functions are given by

$$a_k = c_k \sqrt{\frac{2}{N}} \sum_{n=1}^N g_n \cos \left(\frac{k\pi}{N} \left(n + \frac{1}{2} \right) \right), \quad k = 0, \dots, N-1, \quad (6.10)$$

the coefficients of the discrete sine functions are given by

$$b_k = c_k \sqrt{\frac{2}{N}} \sum_{n=1}^N g_n \sin \left(\frac{k\pi}{N} \left(n + \frac{1}{2} \right) \right), \quad k = 0, \dots, N-1, \quad (6.11)$$

and c_k are the factors to orthonormalize the dictionary

$$c_k = \begin{cases} 1/\sqrt{2} & \text{if } k = 0 \\ 1 & \text{otherwise} \end{cases} \quad (6.12)$$

6.1.4 Choice of regularization parameter and decreasing strategy

Being a regularized estimator, the estimates computed by the MCA algorithm depend on the regularization parameter λ . More degrees of freedom are effectively

¹By considering sine terms, instead of simply using the DCT, it is possible to model sinusoidal signals with arbitrary phase at zero time since $A \cos \theta + B \sin \theta = R \cos(\theta + \phi)$, where $R = \sqrt{A^2 + B^2}$ and $\phi = \tan^{-1}(B/A)$.

used to fit the data as λ decreases to zero. Similar to Sparse PFM, the Akaike Information Criterion (AIC), the Minimum Description Length (MDL) that is equivalent to the Bayesian Information Criterion (BIC), and the universal threshold (UT) and lower universal threshold (LUT) were investigated in Multicomponent PFM.

The UT and LUT values of λ were computed using the maximum absolute deviance (MAD) of the wavelet coefficients (Daubechies wavelets with 2 vanishing moments) at the finest scale level as estimate the standard deviation of the noise, $\hat{\sigma}_{MAD}$ [101].

For AIC and MDL we propose approximating the effective degrees of freedom used to obtain the solution for a given λ as the sum of degrees of freedom to fit the baselines and the degrees of freedom to fit the haemodynamic component [384]

$$\hat{df}_\lambda = \|\alpha_\lambda\|_0 + \text{Tr} \left(\mathbf{H}_\lambda (\mathbf{H}_\lambda^T \mathbf{H}_\lambda)^{-1} \mathbf{H}_\lambda^T \right), \quad (6.13)$$

where $\|\alpha\|_0$ is the number of nonzero (L_0 -norm) coefficients in α_λ , and the matrix \mathbf{H}_λ includes those columns of \mathbf{H} with nonzero coefficients in \mathbf{s}_λ . This approximation is based on the effective degrees of freedom of the LASSO estimates given by (6.6) and (6.7), and it becomes a reasonable approximation provided that the mutual coherence between the dictionaries \mathbf{H} and Φ is negligible [116].

Another important issue in the MCA algorithm is the strategy used to update the regularization parameter since it controls the computational cost of the algorithm and the accuracy of the estimates. For semiparametric PFM, we investigated an exponential strategy [330, 331], where $\lambda^{(n+1)} = \delta \lambda^{(n)}$ and the update factor is given by

$$\delta = \left(\frac{\lambda_{max}}{\lambda_{min}} \right)^{1/(1-I_{max})} \quad (6.14)$$

Choosing a very low number of iterations I_{max} leads to a bad separation of the components. whereas the computational cost of the algorithm considerably increases if I_{max} is too large [43].

6.1.5 Debiasing

Since MCA algorithm is a L_1 -norm regularization, the nonzero coefficients of the selected estimates \mathbf{s}_λ and α_λ tend to be underestimated due to the shrinkage to zero in the search for sparse signals. Consequently, performing a debiasing procedure is recommended to obtain less biased estimates of the morphological components. Two different linear models were investigated for debiasing. The first model comprised

all basis functions of the dictionaries \mathbf{H} and Φ with nonzero MCA coefficients in \mathbf{s} and α , denoted as \mathbf{H}_λ and Φ_λ , respectively. Therefore, the first debiasing linear model to fit to the fMRI voxel time series was:

$$\mathbf{y} = \underbrace{[\mathbf{H}_\lambda \ \Phi_\lambda]}_{\mathbf{X}_1} \beta_1 + \epsilon, \quad (6.15)$$

where $\beta_1 = [\mathbf{s}_\lambda^T \ \alpha_\lambda^T]^T$ includes the subsets of nonzero coefficients in \mathbf{s}_λ and α_λ .

In addition, a second model was also considered:

$$\mathbf{y} = \underbrace{[\mathbf{H}_\lambda \ \mathbf{g}_\lambda]}_{\mathbf{X}_2} \beta_2 + \epsilon, \quad (6.16)$$

where $\mathbf{g}_\lambda = \Phi \alpha_\lambda$ is the estimate of the morphological component and $\beta_2 = [\mathbf{s}_\lambda^T \ \beta_g]^T$ includes the subset of nonzero coefficients in \mathbf{s}_λ and β_g is an amplitude factor that scales the estimate of \mathbf{g}_λ . This second strategy with model \mathbf{X}_2 is more appropriate when the selected λ causes that the representation of the baseline fluctuations stops being sparse, i.e. a large number of basis functions are included to describe the baseline trend, so that the fMRI signal is not overfitted. On the contrary, the complete model \mathbf{X}_1 is more appropriate when sparser representations are favoured with higher values of λ . In practice, the number of nonzero coefficients in α_λ was considered to compute the effective degrees of freedom in (6.13) even though the model \mathbf{X}_2 only employs one degree of freedom to model the baselines. Finally, the least squares estimate of β_i , for $i = 1$ or 2 , is computed as $\beta_i = (\mathbf{X}_i^T \mathbf{X}_i)^{-1} \mathbf{X}_i^T \mathbf{y}$.

6.2 Methods

The properties and performance of Multicomponent PFM were investigated with simulated and experimental fMRI data. The two gamma variate canonical HRF with SPM parameters was used to define the haemodynamic dictionary. The number of iterations of the MCA algorithm was $I_{max} = 50$, and $\lambda_{min} = 0.5 \hat{\sigma}_{MAD}$ is chosen in order to stop the convergence of the MCA algorithm earlier.

6.2.1 Simulated fMRI data

Multicomponent PFM was initially evaluated with the same simulated fMRI data that was generated to investigate Sparse PFM (see Methods section in Chapter 5). In brief, $N_b = 1000$ fMRI voxel time series were generated for each combination of

number of events (0, 2, 4, 6, 8 and 10 events of duration 2s), HRF shape (six different canonical HRF shapes with varying time-to-peak parameter of 3, 4, 5, 6, 7 and 8 s) and temporal signal to noise ratio (tSNR) between 30 and 80. Since the HRF used in the model has a time-to-peak of 5 s, it was possible to assess the performance of the Multicomponent PFM in HRF mismatch scenarios due to variability in the actual HRF. Two different noise models were investigated: additive white gaussian noise (AWGN), and AWGN plus two sinusoidal components with 4 harmonics each (8 sinusoids in total), which aim to simulate physiological cardiac and respiratory fluctuations (see Eq. (5.17)). The frequencies of the main frequency and harmonics were randomly generated as $f_{.,i} \sim N(if., 0.04)$, with $f_r = 0.3$ Hz [38] and $f_c = 1.1$ Hz [320]. A power-of-2 decrease in the amplitude of harmonics was also simulated [144]. The ratio between the sinusoidal and AWGN components of the noise was modelled following [343]. The fMRI time series were 256 s in duration and were sampled at TR 2 s, resulting in $N=128$ time points. Note that the sinusoidal components of the signal are undersampled at this simulated TR since the Nyquist frequency is 0.25 Hz.

Similar to the simulations performed for Sparse PFM in Chapter 5, we measured the sensitivity and specificity in detecting the simulated activation events over time with (5.19) and (5.20). On the other hand, we also measured the mean square error of the estimate of the BOLD component (MSE_s) computed with (5.18), and the mean square error of the sinusoidal fluctuations (MSE_g) which was defined as

$$MSE_g = \frac{1}{N_b} \sum_{i=1}^{N_b} \|\mathbf{g}_i - \Phi \hat{\boldsymbol{\alpha}}_i\|_2^2, \quad (6.17)$$

in order to examine the accuracy of the computed estimates of the BOLD signal and the sinusoidal trend and evaluate the robustness of the method to recover the BOLD and sinusoidal components of the signal.

Each voxel time series was analyzed with Multicomponent PFM considering the four methods to select the regularization parameter (AIC, MDL, UT and LUT), and the two debiasing strategies described in (6.15) and (6.16). In addition, we compared the performance of the method with the GLM analysis in Chapter 5. The GLM comprised one regressor per activation event, which was created as the convolution of the Canonical HRF with the simulated stimulus function of the event.

6.2.2 Experimental fMRI data

We illustrated the principle of Multicomponent PFM with the datasets acquired at TR 2 s for the visuomotor paradigm used in Chapters 4 and 5. More details about the experiment are given in Chapter 4.

For this study, the original fMRI datasets were simply corrected for motion and a linear trend with AFNI (NIMH/NIH) [82]. Afterwards, the datasets were analyzed with Semiparametric PFM where the debiasing was done with \mathbf{X}_2 (6.16) since this model has lower computational cost than the model \mathbf{X}_1 and little difference was observed in the simulation results between both debiasing models with respect to the estimation of the haemodynamic component of the signal. Total, positive and negative Activation Time Series (ATS) counting the number of voxels with nonzero coefficients at each time point were also created for each λ -selection approach.

Spearman's nonparametric rank correlation coefficients [171] between the ATS and the EMG signals from the left extensor (LE), right extensor (RE) and right flexor (RF) were computed. After decimation to the same length as the fMRI data, the EMG signals were thresholded in order to indicate only significant EMG activity (amplitude of the decimated signal is larger than 4 times the standard deviation of the signal).

Multicomponent PFM Activation Maps were created for the finger tapping events. The maps displayed the maximum absolute value (maximum or minimum) of the coefficients of \mathbf{s} during the time points defining the Activation Event. A minimum cluster size of 5 neighbouring voxels was employed in order to reduce isolated activations that are likely to be false positives. These maps were compared with the results of the GLM analysis performed in Chapter 5 in terms of overlapping voxels ($\text{PFM} \cap \text{GLM}$) and its ratio to the number of voxels detected with Multicomponent PFM and GLM.

Finally, since we aimed to represent the baseline fluctuations with a morphological dictionary composed of sine and cosine basis functions, we used RETROICOR [144] as a benchmark to evaluate the ability of the method to extract sinusoidal fluctuations of physiological origin. RETROICOR was applied to the motion corrected data, prior to linear detrending. The comparison between both methods was done in the frequency domain and we computed the power spectral density of the signal removed by RETROICOR and the baseline component estimated by Multicomponent PFM.

6.3 Results

6.3.1 Simulated data

Figure 6.1 illustrates the idea behind Multicomponent PFM in a simulated fMRI time series with 5 events of duration 2 s (4 positive events and 1 negative event). For this example, the debiasing model \mathbf{X}_2 was used. The simulated HRF to generate the BOLD responses had a time to peak of 7 s, whereas the HRF of the model has a time to peak of 5 s. The simulated temporal signal-to-noise ratio (tSNR) and contrast-to-noise ratio (CNR) of the time series was 50 and 3, respectively. It can be observed that the MCA algorithm used in Multicomponent PFM is able to separate the BOLD and sinusoidal components of the signal if the UT and LUT approaches are used to select λ . Setting λ using UT produces a false negative at 98 s, whereas a false positive is detected with LUT at 196 s. In contrast, setting λ according to a model selection criteria results in overfitting in the case of AIC and choosing an overconservative null estimates in case of MDL. In the Fourier domain, it can be seen that the technique can detect the main frequency components and the harmonics of the simulated respiratory and cardiac signals with the UT and LUT criteria, whereas the AIC criteria starts fitting the spectral components of the random white noise.

Figure 6.2 depicts the receiver operating characteristic (ROC) curves showing the specificity and sensitivity values achieved by Multicomponent PFM for the debiasing model \mathbf{X}_2 in Eq. (6.16). Differences in sensitivity and specificity to detect the BOLD responses between both debiasing models were negligible. The curves correspond to a scenario where perfect knowledge of the HRF is available (No HRF mismatch) and when the simulated HRF is the canonical HRF with time to peak of 8 s in contrast to the HRF of the model which has a time to peak of 5 s. It can be seen that Multicomponent PFM achieved very high specificity values (less than 5% false positive rate) with all the investigated criteria in case of perfect knowledge of the HRF, and less than approximately 10 % false positive rate with HRF mismatch. In general, we observed that specificity values are higher than those obtained with Sparse PFM in the same simulations (see Figure 5.4 in Chapter 5), demonstrating the benefit of modelling the baseline fluctuations. In terms of sensitivity, we observed that higher values for the AIC and LUT than in Sparse PFM, whereas similar results were achieved with UT. The most relevant observation of the ROC analysis with Multicomponent PFM is the good operation of AIC, which is now

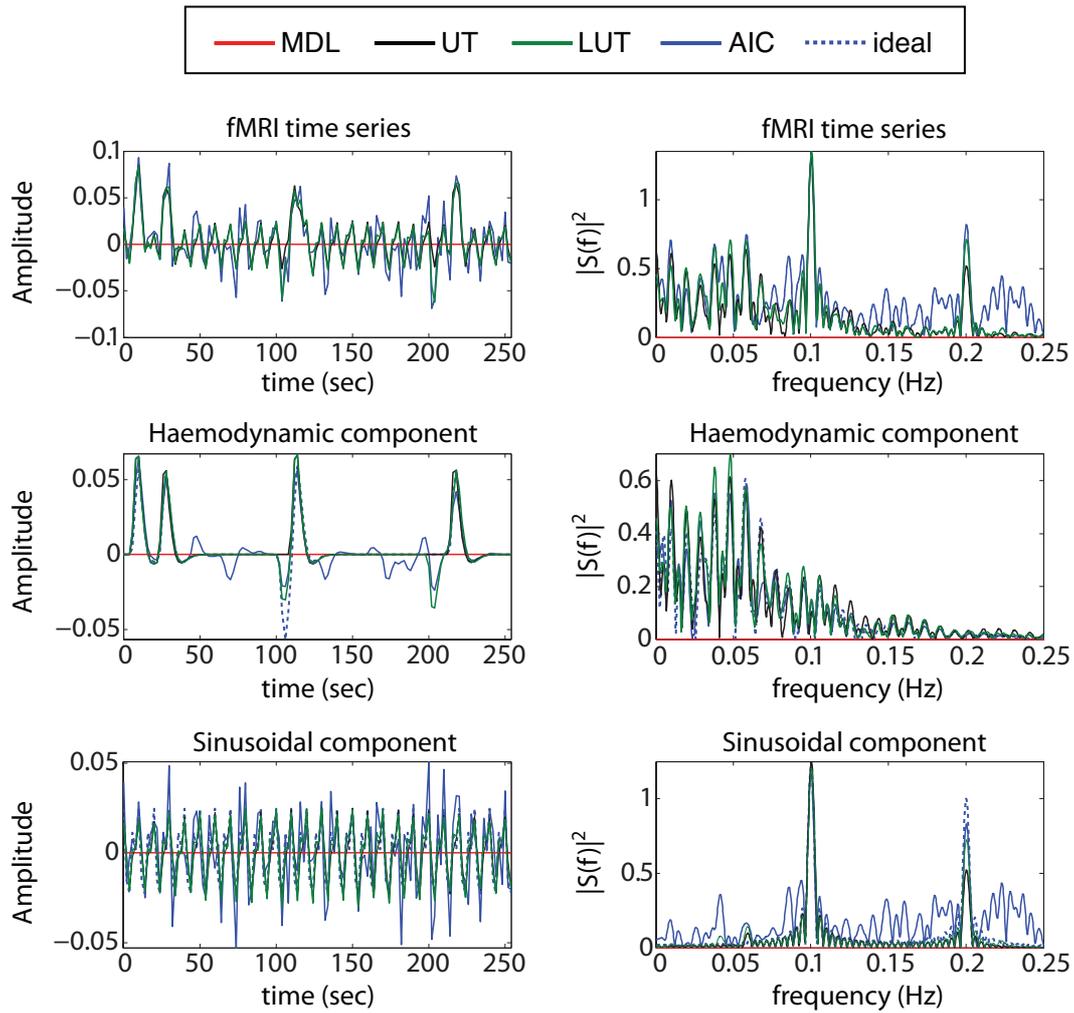


Figure 6.1: Example of the operation of Multicomponent PFM in a simulated fMRI time series ($t\text{SNR} = 50$, $\text{CNR} = 3$, $\text{TR} = 2$ s). The left figures plot the original simulated time series (top) and the haemodynamic (middle) and sinusoidal (bottom) components estimated by the MCA algorithm, whereas the right figures plot the corresponding spectra of the fMRI time series and each component.

comparable to the UT and LUT approaches. This contrasts to the unsatisfactory results obtained for Sparse PFM when AIC was used to select the regularization parameter λ . Furthermore, using the MDL criterion to select λ in Multicomponent PFM resulted in a considerable reduction in sensitivity of the method where the number of simulated events increased to 6 and 10.

Figures 6.3 and 6.5 show the Mean Square Error of the estimate of the BOLD component (MSE_s) as a function of the number of events and $t\text{SNR}$, respectively, while Figures 6.4 and 6.6 plot the mean square error of the estimate of the sinusoidal trends (MSE_g) as a function of the number of events and $t\text{SNR}$, respectively. These figures demonstrate that using the MDL criteria in Multicomponent PFM

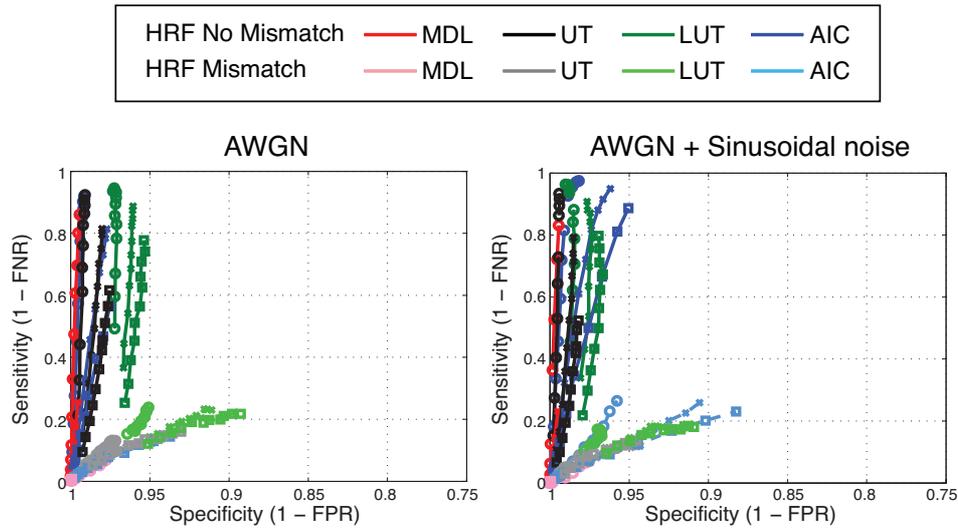


Figure 6.2: Receiver Operating Curves (Sensitivity vs. Specificity) of Multicomponent PFM for the debiasing model \mathbf{X}_2 in Eq. (6.16). The model HRF is the Canonical HRF with standard SPM parameters. No mismatch scenario: Simulated HRF is equal to the model HRF with a time-to-peak of 5 s; HRF Mismatch scenario: Simulated HRF has a time to peak of 8 s. Line markers indicate number of activation events: 2 (\circ), 6 (\times) and 10 (\square). Points with decreasing sensitivity in each line correspond to tSNR decreasing from 80 to 30.

is not recommended since it computes inaccurate estimates of the haemodynamic and sinusoidal component when these are present in the fMRI signal. The poor operation of MDL is due to the high values of λ chosen with this criterion which causes excessive shrinkage of the coefficients to zero. This effect in turn becomes beneficial when the sinusoidal fluctuations must not be estimated and explains the low MSE_g obtained with MDL when only AWGN is simulated in Figures 6.4 and 6.6. In detecting the BOLD responses, the MDL criterion would only give satisfactory operation at high tSNR conditions and with very low numbers of events. This contrasts with the behaviour shown by MDL for Sparse PFM where it demonstrated excellent operation. This suggests that the penalization applied by MDL (equal to $\log N$) is too restrictive due to the rapid increase in the effective degrees of freedom used to fit the time series when modelling the baselines and the haemodynamic component. In general, selecting λ based on the UT and LUT outperformed the results obtained with model selection criteria (AIC and MDL). Selecting λ with LUT exhibits improved estimation properties when the number of events is large due to the higher sensitivity of the method and lower degree of shrinkage. However, this characteristic becomes a drawback when no events exist the fMRI signal (2-3% higher false positive rate than UT) and setting λ to the UT provides better control for detecting false positives. Furthermore, regularizing with the L_1 -norm provides

certain robustness of the method against HRF mismatches in that the technique is able to provide reasonable estimates of the haemodynamic component of the signal despite the decrease in sensitivity observed in the ROC curves. Again, the type of debiasing did not affect the estimate of the haemodynamic component (see Figures 6.3 and 6.5). In contrast, the type of debiasing becomes relevant if one is interested in estimating the baseline fluctuations of the signal. According to the simulation results, debiasing using \mathbf{X}_2 attained better estimates than debiasing with \mathbf{X}_1 where each nonzero coefficient describing the haemodynamic and sinusoidal components are treated independently. Furthermore, Multicomponent PFM outperforms the GLM analysis in case of HRF mismatch at high SNR conditions and with lower simulated activation events and the technique asymptotically converges to the optimal estimates provided by GLM in case of no HRF mismatch as the tSNR of the signal increases. In comparison with Sparse PFM, we observed that Multicomponent PFM improved the estimate of the haemodynamic component, i.e. lower MSE_s , when λ was selected based on the LUT, whereas similar curves were obtained for UT.

6.3.2 Experimental data

The results in experimental data demonstrated that Multicomponent PFM can decouple sinusoidal fluctuations of the fMRI signal, including physiological fluctuations, and single trial BOLD responses without prior timing information of the events. Figure 6.7 depicts the Multicomponent PFM decomposition of an fMRI time series located in the motor cortex, fitting separately the haemodynamic component of the signal and the baseline fluctuations. Note that the technique is able to identify the main sinusoidal trends of the signal (strong peaks at 0.005 Hz, 0.014 Hz and 0.23 Hz), along with the 4 finger tapping trials without using prior information about the events or physiological information

Figure 6.8 plots the Activation Time Series (ATS) obtained with Multicomponent PFM using UT, which was the criteria with which a smaller number of activations were found. It can be seen that all visually cued and self-paced finger tapping events were detected in all datasets. The ATS computed by Multicomponent PFM using UT showed fewer voxels deemed active than the corresponding ATS computed with Sparse PFM (note that the ATS presented in Chapter 5 were computed using MDL). As shown in Table 6.1, the Spearman's correlation coefficients between the ATS and the EMG signals showed no significant correlation in the LE of Subject B, the RE and RF of Subject C and the EMG signals of Subject D, even the Activation Events

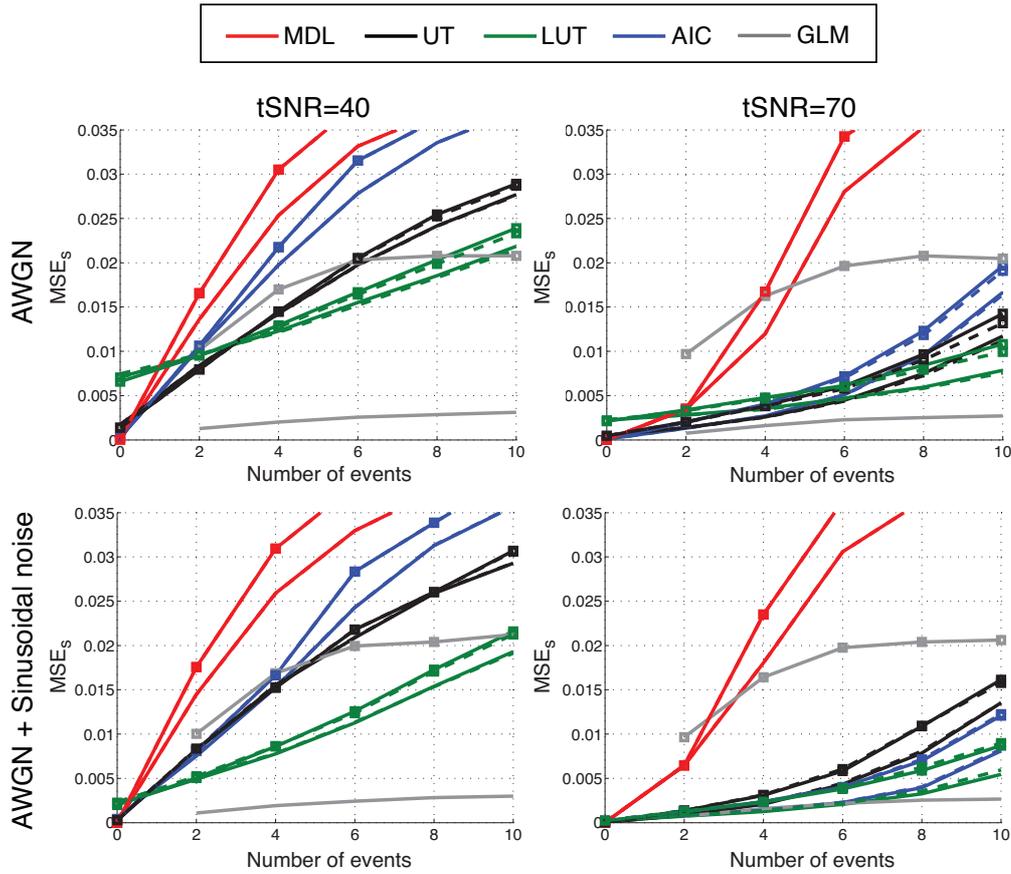


Figure 6.3: MSE between the simulated and the estimated stimulus signal (MSE_s) as a function of the number of activation events. Solid lines: Debiasing model (6.15); Dash lines: Debiasing model (6.16). Square markers (\square) correspond to HRF Mismatch (Canonical HRF with time-to-peak of 8 s), whereas lines without markers correspond to no HRF mismatch (Canonical HRF with time-to-peak of 5 s).

for this datasets can be clearly identified in the ATS. This divergence is related to the larger amount of activations found during the periods of rest and may not be related to hand movements and therefore is not captured by the EMG recordings.

Figure 6.9 depicts the Multicomponent PFM maps (amplitude of coefficients) and GLM maps (F-test, $Q < 0.005$, FDR corrected) for the 4 finger tapping events of the datasets of subject D. The Multicomponent PFM maps show the maximum absolute value of the coefficients during the Activation Event, whereas the GLM maps plot the amplitude of the coefficient corresponding to the canonical HRF regressor. The figure demonstrates strong similarity in the spatial distribution of the activations between both techniques, even though quantitatively we observed that 36.7 % of the activations detected with Multicomponent PFM were also present in GLM, and 68.7 % of the GLM activations were also present in Multicomponent PFM. It can be observed that cortical responses associated with single-trial finger

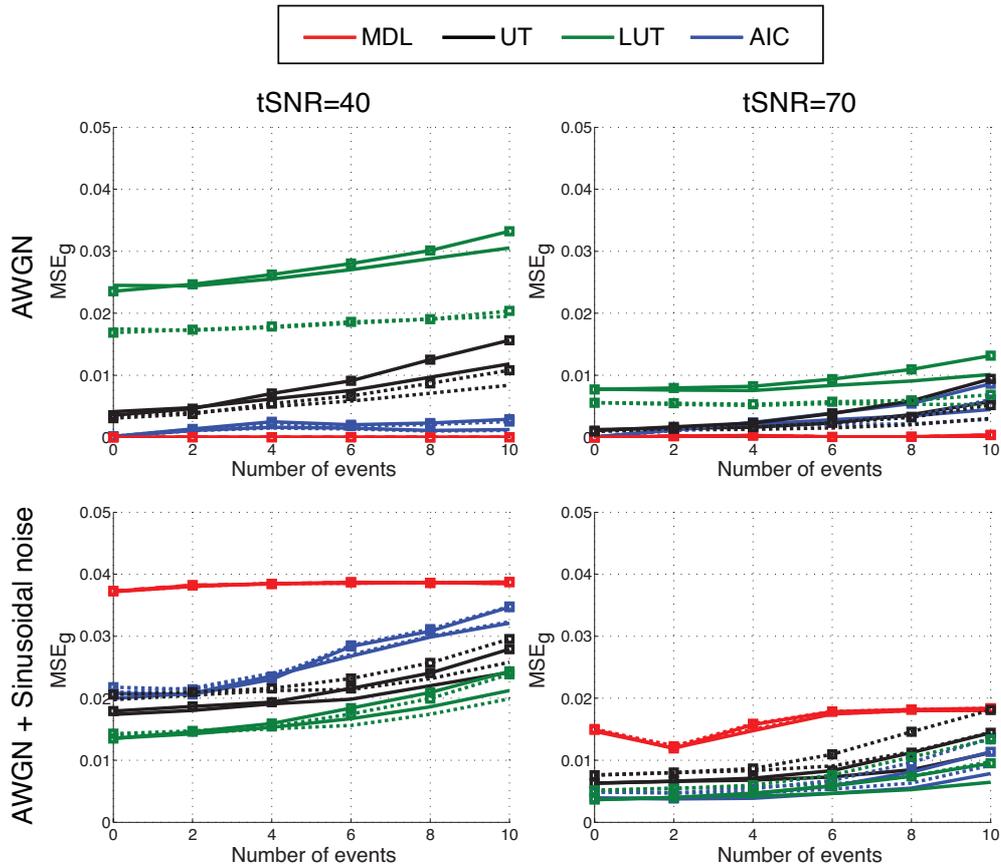


Figure 6.4: MSE between the simulated and estimated sinusoidal trend (MSE_g) as a function of the number of activation events. Solid lines: Debiasing model (6.15); Dash lines: Debiasing model (6.16). Square markers (\square) correspond to HRF Mismatch (Canonical HRF with time-to-peak of 8 s), whereas lines without markers correspond to no HRF mismatch (Canonical HRF with time-to-peak of 5 s).

tapping events were found in areas of the bilateral primary motor, bilateral primary somatosensory, bilateral superior parietal cortices, intraparietal sulcus and visual cortex. The technique also displayed a higher number of active voxels, i.e. higher sensitivity, in the events detected during periods of rest and these activations also showed close spatial concordance with those detected with PFM and Sparse PFM.

Figures 6.11 and 6.10 depict the maps of the Fourier coefficients at the respiratory (0.23Hz) and cardiac (0.014Hz) frequencies for Subject D. These frequencies were selected based on the spectrum shown in Figure 6.7. It can be seen that the spatial distribution of baseline signal estimated with Multicomponent PFM showed large agreement with those of RETROICOR. In general, we observed that RETROICOR removed more frequency components than Multicomponent PFM. It can be seen that the maps computed with lower regularization parameters (AIC and MDL) criteria showed a higher resemblance with the RETROICOR maps. The regions

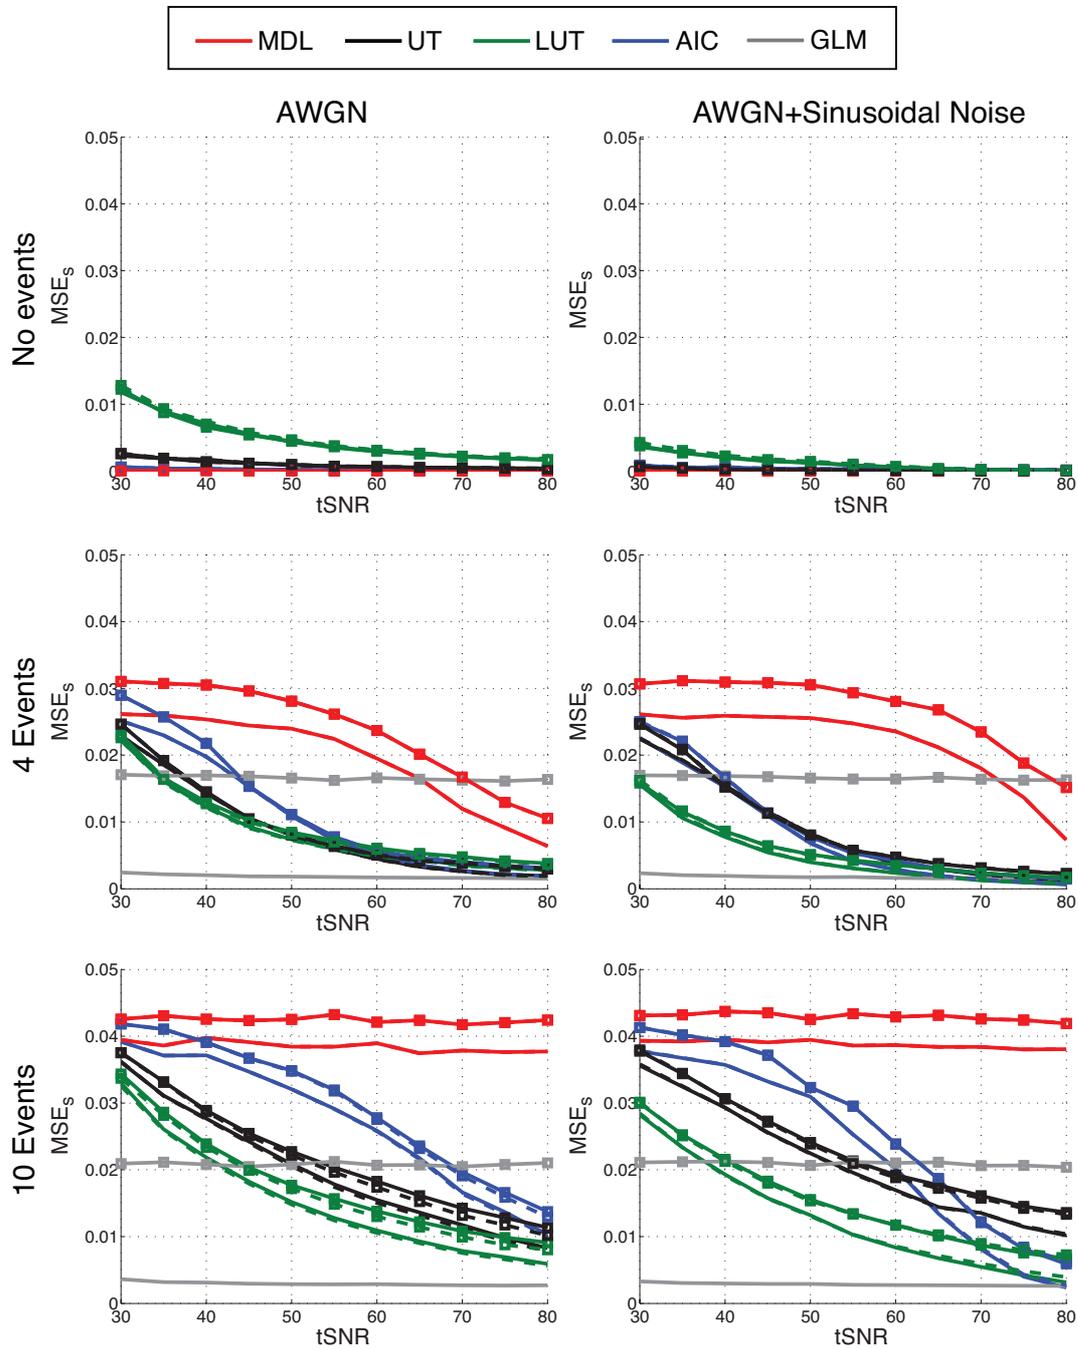


Figure 6.5: MSE between the simulated and the estimated stimulus signal (MSE_s) as a function of tSNR. Solid lines: Debiasing model (6.15); Dash lines: Debiasing model (6.16). Square markers (\square) correspond to HRF Mismatch (Canonical HRF with time-to-peak of 8 s), whereas lines without markers correspond to no HRF mismatch (Canonical HRF with time-to-peak of 5 s).

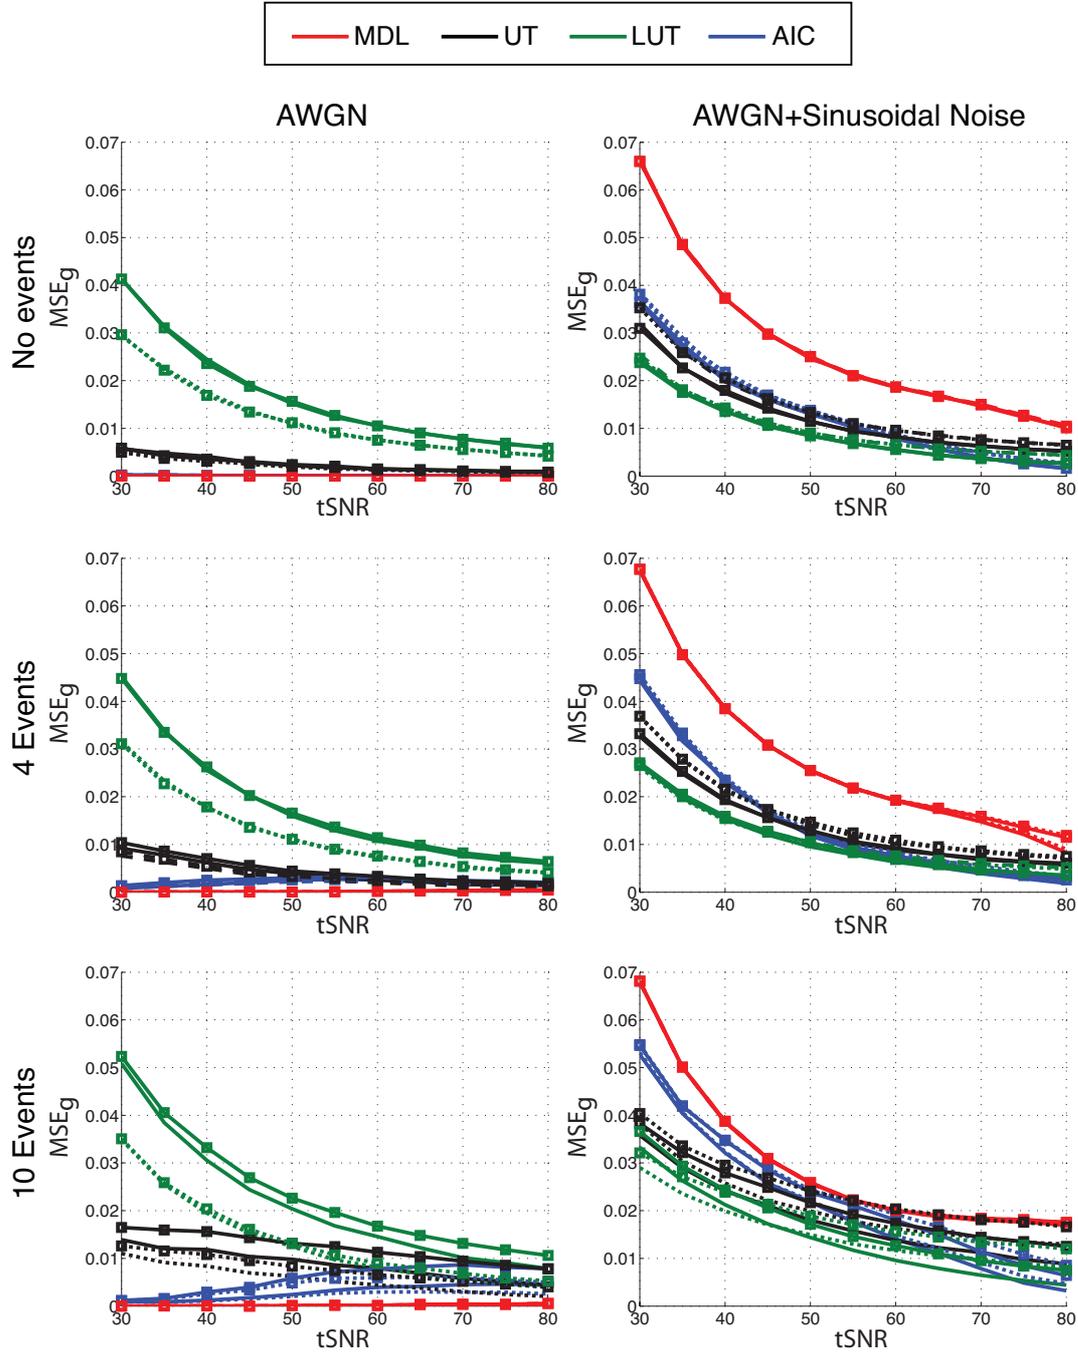


Figure 6.6: MSE between the simulated and estimated sinusoidal trend (MSE_g) as a function of $tSNR$. Solid lines: Debiasing model (6.15); Dash lines: Debiasing model (6.16). Square markers (\square) correspond to HRF Mismatch (Canonical HRF with time-to-peak of 8 s), whereas lines without markers correspond to no HRF mismatch (Canonical HRF with time-to-peak of 5 s).

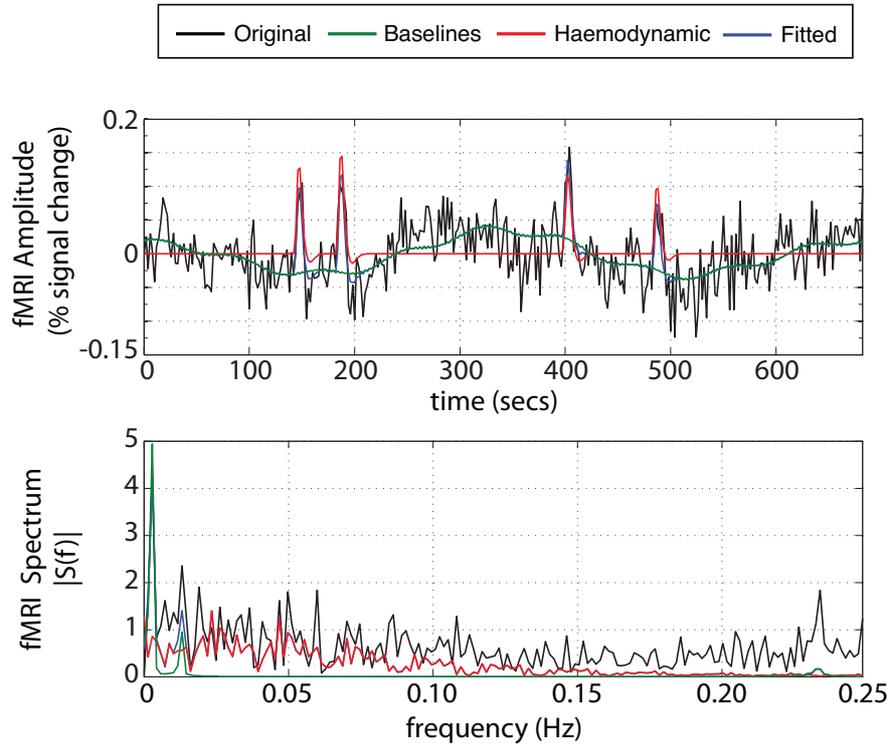


Figure 6.7: Multicomponent PFM method using the UT applied to a fMRI time series located in the motor cortex. The top figure plots the original fMRI data (just after regressing out for a linear trend) in black, and the estimated haemodynamic component (\mathbf{H}_s) in red, baseline fluctuations (\mathbf{g}) in green and fitted signal ($\mathbf{H}_s + \mathbf{g}$) in blue. The bottom figure plots the corresponding spectra of the signal. It can be observed that the method is able to identify the four finger tapping events without prior timing information, and identifies the main sinusoidal frequencies of the baselines (strong peaks near 0.005 Hz, 0.014 Hz and 0.23 Hz).

with higher spectral amplitude in the cardiac component were localized in the edges of the brain such as near sulci, or major vessels such as the sagittal sinus or the middle cerebral artery. Relevant coefficients from respiratory-induced fluctuations were more scattered across the brain, and also showed overlapping regions with cardiac fluctuations. Similar results were obtained for the rest of the datasets.

6.4 Discussion

These results illustrate that Multicomponent PFM can decompose the fMRI voxel time series into a haemodynamic component without prior information about timing of the events and sinusoidal baseline fluctuations without specifying the frequencies to remove. In consequence, our results have illustrated that the technique is able to extract the main cardiac and respiratory components of the fMRI signal and showed good correspondance with RETROICOR [144].

	Left Extensor Spearman's	Right Extensor Spearman's	Right Flexor Spearman's
Subject A	0.263*	0.286*	0.233*
Subject B	0.054	0.172*	0.260*
Subject C	0.184*	0.173	0.153
Subject D	0.245*	0.217*	0.201*
Subject F	-0.064	0.062	0.149

Table 6.1: Spearman's correlation coefficients between the decimated EMG signals and the ATS computed with Multicomponent PFM with λ chosen as the Universal Threshold (UT). The asterisks indicate that the correlation coefficient is statistically significant (p -value < 0.001).

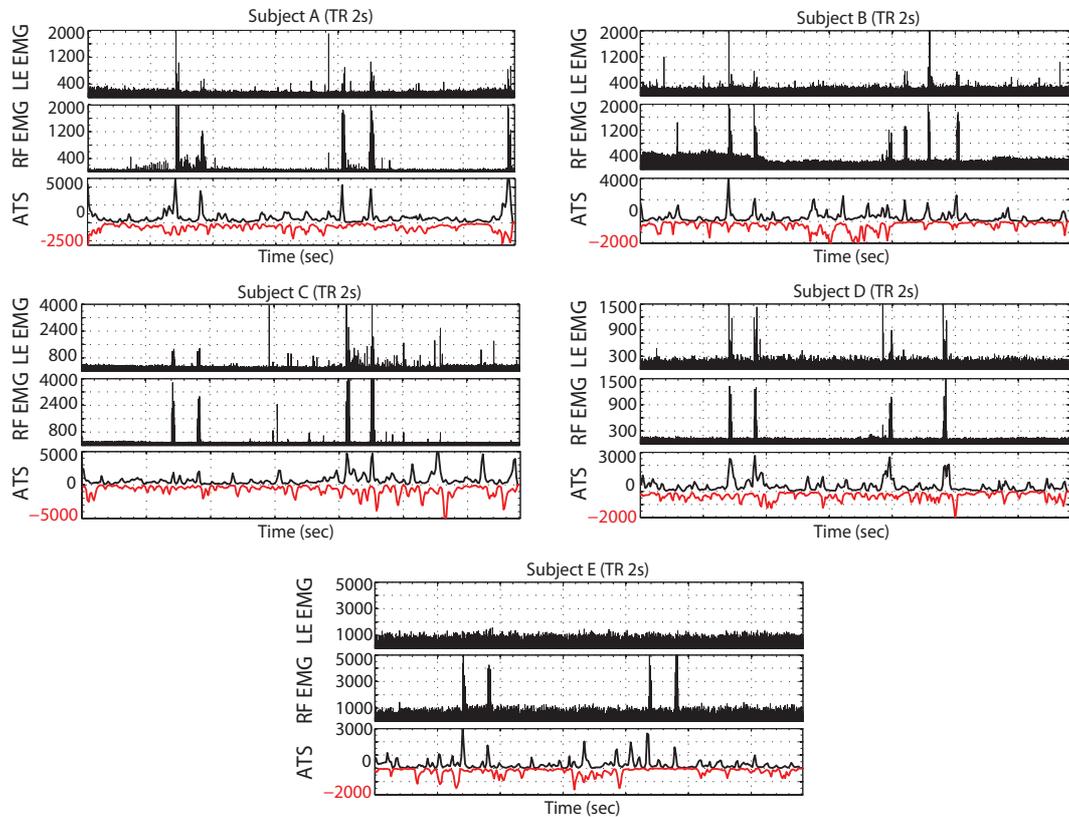


Figure 6.8: Activation Time Series (ATS) computed with Multicomponent PFM (UT criteria) (bottom), left extensor (LE) (top) and right flexor (middle) EMG signals. In order to facilitate the interpretation of the result, the ATS differentiates between positive activations (black, positive y-axis) and negative activations (red, negative y-axis).

Methodological issues and simulations

Multicomponent PFM is based on partially linear model formulation of the fMRI voxel time series and relies on morphological component analysis (MCA) to separate the signal into two morphological components. The present work extends upon previous semiparametric approaches for fMRI data analysis [116, 231, 251]. This

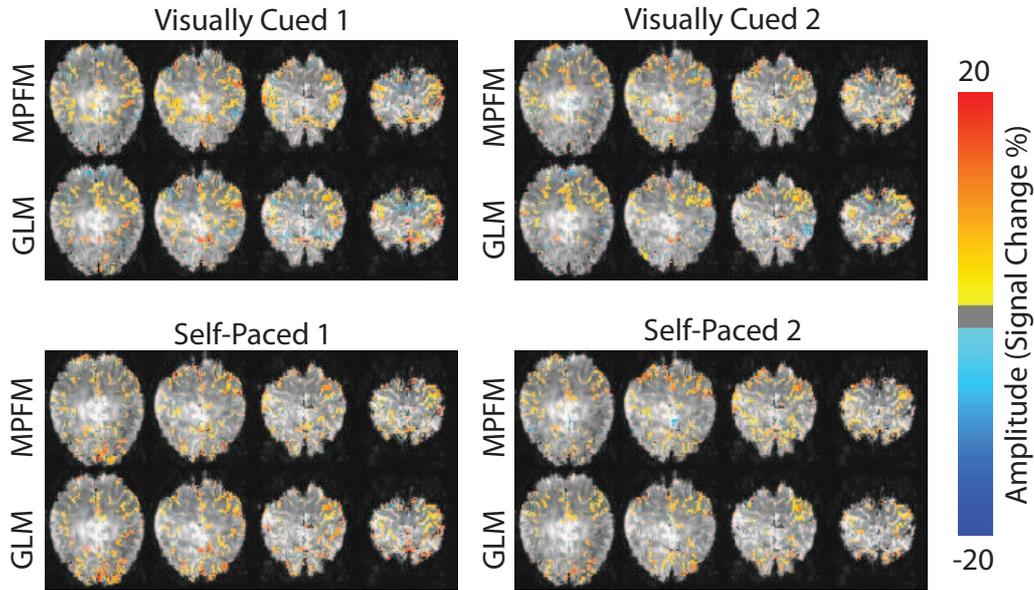


Figure 6.9: Comparison between the Multicomponent PFM maps and the GLM maps ($Q < 0.05$, FDR corrected) for the 4 finger tapping events of the datasets of subject D. The Multicomponent PFM maps show the maximum absolute value of the coefficients during the Activation Event, whereas the GLM maps plot the amplitude of the coefficient corresponding to the canonical HRF regressor.

previous work proposed semiparametric generalized linear models where the parametric part described the effect of experimental stimuli and thus required information about the paradigm. In these approaches the baseline drift was also modelled as a linear combination of basis functions from a dictionary. Large scale wavelets were used in [231, 251], whereas our approach does not specify the frequency or scale of the baseline functions, similar to the general formulation developed in [116]. To fit the semiparametric model, a Maximum Likelihood estimator is employed in [251], sparse Bayesian learning in [231] and penalized Partially Linear Model estimators or the MCA algorithm so that the experimental effects are estimated via least squares estimation [116]. Fundamentally, the main contribution of our work is to extend the semiparametric models into a paradigm free scenario where no specification of the experimental effects is required. This becomes feasible due to the sparse estimation of the BOLD responses with the MCA algorithm.

Modelling the baseline fluctuations with a dictionary comprising cosine and sine basis functions was proven useful in this study, but it would not be appropriate for periodic experimental paradigms [129]. In that case, caution should be taken so that the removed baseline fluctuations are not synchronous with task activations similar to the recommendations for other noise or drift correction techniques [144, 336]. Due to the stationary of the cosine and sine transforms, the dictionary can only model

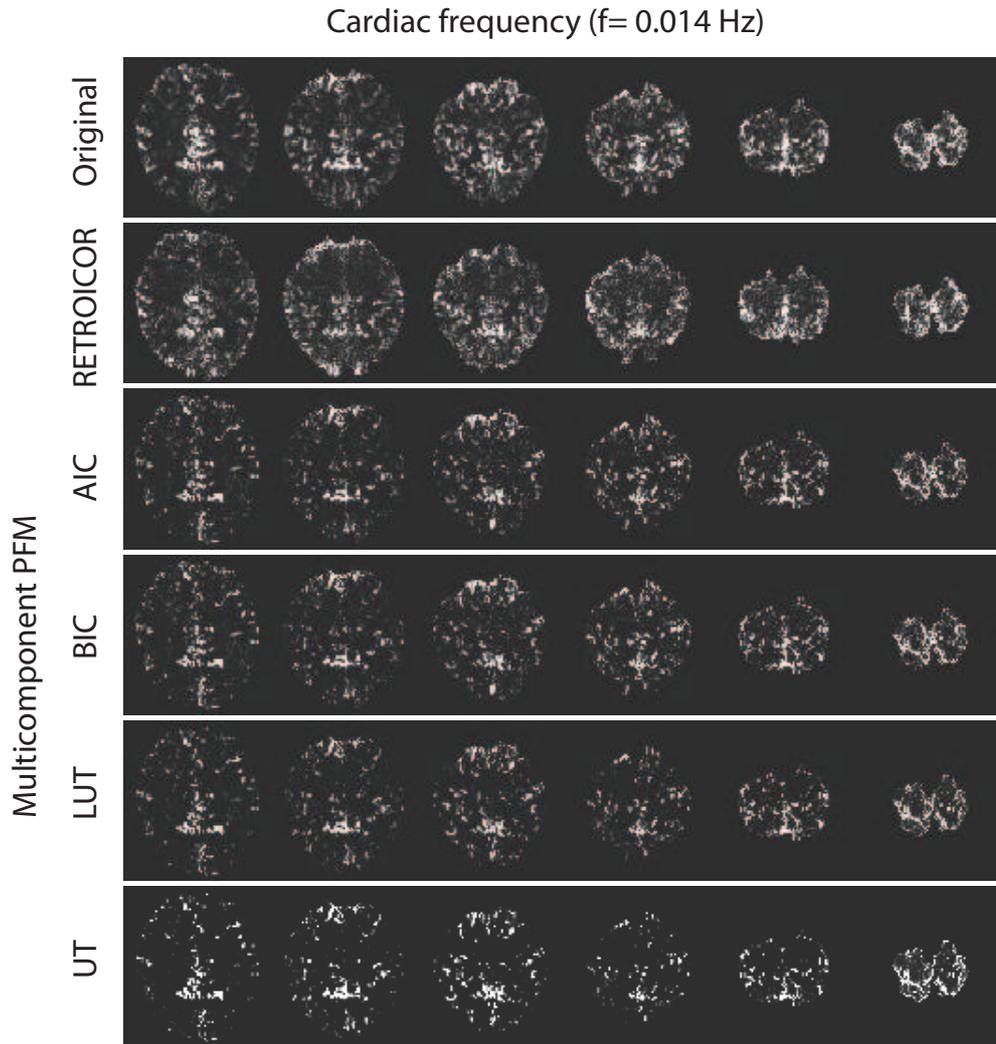


Figure 6.10: Maps of the Fourier coefficient at frequency 0.014 Hz (cardiac) in six different slices of dataset of Subject D. The rows depict maps computed from the original fMRI data, the component removed by RETROICOR and the baselines components estimated with Multicomponent PFM with the AIC, MDL, LUT and UT criteria. In this case, the cardiac-related fluctuations were more localized in the edges of the brain and major vessels.

stationary trends of the baseline. In future work, dictionaries based on wavelets could also be explored to model nonstationary baseline drifts of the signal provided the mutual coherence of the dictionaries is controlled. In that sense, wavelet dictionaries can handle the typical $1/f$ spectrum of the noise observed in fMRI data [231, 251]. Here, this issue was tackled in the experimental analysis by detrending the time series with a linear regressor as suggested in [35].

We found that setting the regularization parameter to UT and LUT outperformed selection based on model selection criteria. In our simulations, the MDL criteria exhibited the worst performance in terms of sensitivity and Mean Square Errors (MSE_s and MSE_g) for Multicomponent PFM. This contrasts with the re-

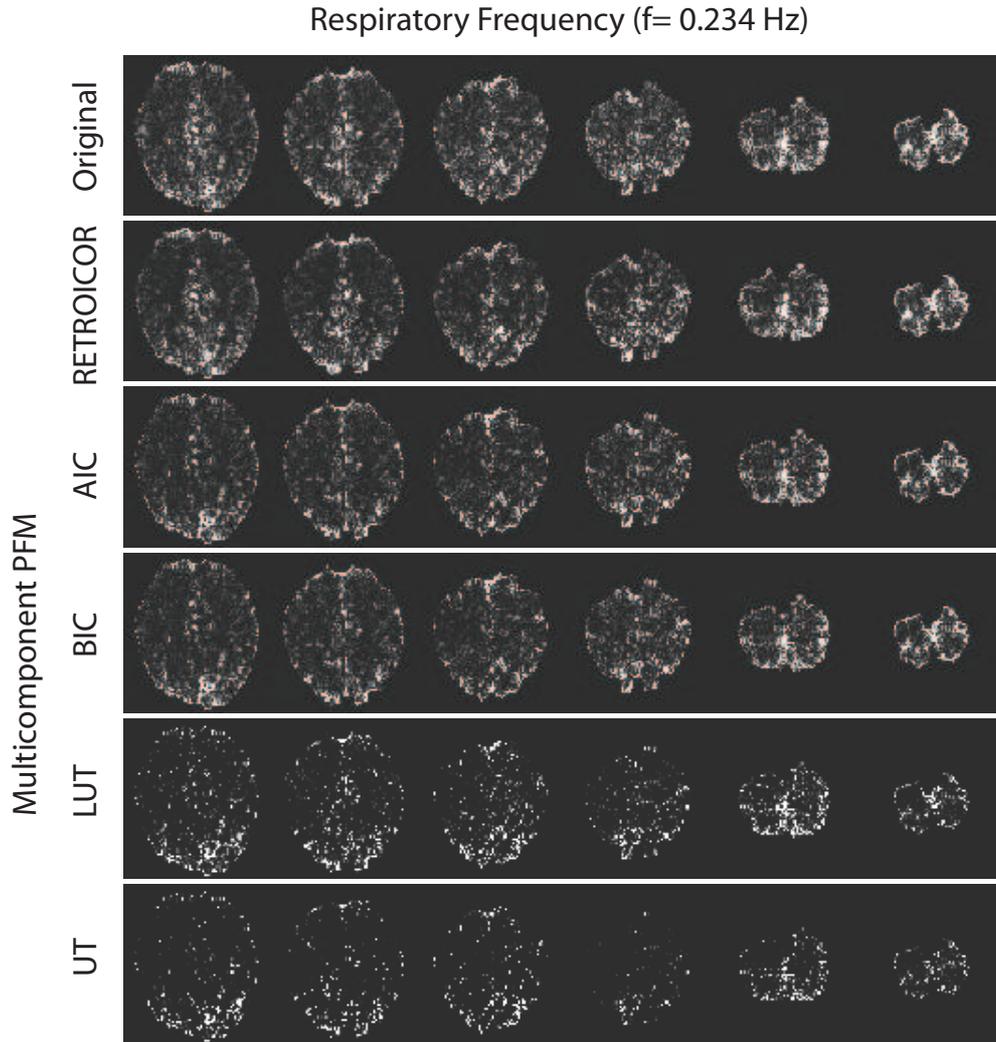


Figure 6.11: Maps of the Fourier coefficient at frequency 0.234 Hz (respiratory) in six different slices of dataset of Subject D. The rows depict maps computed from the original fMRI data, the component removed by RETROICOR and the baselines components estimated with Multicomponent PFM with the AIC, MDL, LUT and UT criteria. In this case, the respiratory-induced fluctuations were more scattered across the brain.

sults for Sparse PFM in which MDL gave satisfactory results. Models with large effective degrees of freedom are penalized more with MDL than with AIC. Consequently, very high values of λ are selected which results in excessive shrinkage of the haemodynamic and DCST coefficients. In agreement with [116], we used the UT in the experimental analysis since it achieves better false positive control in the simulations. Investigating alternative approaches for selecting λ in a regional basis instead of voxelwise will be a topic of future research.

The simulations also demonstrated robustness of the technique to HRF model mismatches owing to the L_1 -norm regularization, similar to the results observed in Chapter 5 for Sparse PFM. It was observed that the main advantage of the MCA

method relative to Sparse PFM is attained at high SNR regimes where the ratio of the morphological sinusoidal component with respect to AWGN is so large that this component can be accurately estimated and then subtracted so that the estimate of the haemodynamic component is improved. Therefore, Multicomponent PFM is especially practical for event related fMRI studies at ultrahigh (7T) MR field and/or acquisition with large voxel sizes where physiological noise is predominant over thermal noise, whereas we recommend using Sparse PFM in high-resolution fMRI studies where the physiological noise of the signal is reduced [44, 343]. According to our simulations, Sparse and Multicomponent PFM techniques are able to detect single trial responses with high specificity at tSNR values as low as 30 or 40, which corresponds to a 3 mm^3 voxel resolution at 7T [343]. Establishing the minimum spatial resolution at which single trial responses can be experimentally detected becomes now a very relevant research question.

No statistical inference of the estimates was provided in this study. Activation Maps and Activation Time Series in Multicomponent PFM were computed without statistical inference and thresholding. This decision is rooted in our simulation results which demonstrated that the technique is able to operate at very low false positive rates. Future work will study the asymptotic properties of the estimator and statistical inference in partially linear models in the context of PFM [116, 383]. Fadili and Bulmore (2005) suggested to use nonparametric inference techniques, such as bootstrapping, to build confidence intervals of the estimates [116]. However, note that using nonparametric inference in Multicomponent PFM would considerably increase the computational cost of the method, becoming an impractical alternative at the current computational speed.

Experimental results

The usefulness of Multicomponent PFM was demonstrated through the analysis of the experimental datasets of a single trial visuomotor paradigm and using the RETROICOR algorithm as a benchmark to evaluate whether Multicomponent PFM is able to remove baseline physiological fluctuations without the need of external physiological monitoring. This can be relevant in clinical cases where physiological monitoring can be difficult or when the jitter of the phase measures obtained from the pulse-oximeter or the ECG is unacceptable [229].

The spatial distribution of the coefficients at the cardiac and respiratory frequencies showed larger correpondance with the spatial frequency maps of the raw data

and the signal removed by RETROICOR. These components were mainly located in GM and CSF voxels, and near major vessels [144, 200, 229, 362], with a more distributed map for the respiratory component [144]. Spatial overlapping between both frequency components might be due to the short TR employed in this study which causes aliasing of respiratory and cardiac harmonics [144, 229]. Higher sampling frequencies will help to distinguish between both components and this has been observed in simulations.

The resemblance between the spatial distribution of the cardiac and respiratory components extracted with RETROICOR and Multicomponent PFM increased as the regularization parameter was reduced (AIC and MDL criteria), whereas using the UT and LUT approaches resulted in sparser maps. However, quantitative comparison of both methods was not feasible in this study due to differences in pre-processing of the signals input to both methods, which will have biased the results. Future work will address a quantitative investigation of the components of the fMRI signal removed by Multicomponent PFM and those removed by RETROICOR and other physiological noise correction techniques, such as respiratory volume regressors from RVT [38, 40] or cardiac-rate regressors [71, 320] in terms of the coefficient of determination, R^2 , and the variability of the time series explained by each signal component [35, 185].

All in all, both observations (larger sensitivity of Multicomponent PFM to detect the single trial responses than Sparse PFM and better reproducibility of the physiological fluctuations with lower λ) calls for the use of two independent regularization parameters for each morphological component. We feel that larger values of λ will be more adequate in fitting the haemodynamic components in order to control false positives. However, considering RETROICOR as a benchmark, lower values for λ seem to be more appropriate in order to further improve the estimation of the baseline physiological components of the signal. Heuristic rules can be formulated to facilitate the search of the optimal regularization parameters in the two-dimensional plane [98, 318].

A visionary fMRI data analysis technique

It might be the case that the most important result from Multicomponent PFM is still to be explored. Had dictionaries been defined for each of potential component

of the fMRI signal, the model used in Multicomponent PFM could be extended as

$$\begin{aligned} \mathbf{y} &= \mathbf{H}\mathbf{s} + \sum_{i=1}^K \mathbf{g}_i + \boldsymbol{\epsilon} \\ &= \mathbf{H}\mathbf{s} + \sum_{i=1}^K \boldsymbol{\Phi}_i \boldsymbol{\alpha}_i + \boldsymbol{\epsilon}, \end{aligned} \quad (6.18)$$

where each morphological component \mathbf{g}_i can be represented as a linear combination of basis functions from a dictionary $\boldsymbol{\Phi}_i \boldsymbol{\alpha}_i$. Subsequently, a block coordinate relaxation algorithm, similar to the MCA procedure, could be used to recursively decompose the fMRI signal into the multiple morphological components. In that scenario spatio-temporal models of the components could be proposed in order to make use of the spatio-temporal smoothness of fMRI data. Ultimately, morphological decomposition of the fMRI signal could be used as a paradigm-free model-based alternative to other multivariate fMRI analysis techniques which decompose the signal according to orthogonality criteria in principal component analysis (PCA) [8], or independence [65] or sparsity [89] criteria in independence component analysis (ICA) or dictionary learning decomposition [208].

6.5 Conclusion

Multicomponent PFM has been presented in this chapter based on the MCA algorithm. This technique can automatically detect BOLD fMRI responses without prior information about the paradigm and also characterize and correct baseline trends due to physiological cardiac and respiratory fluctuations without the need for physiological monitoring. Simulation results have demonstrated higher sensitivity and specificity than Sparse PFM. Our experimental results have illustrated the paradigm-free detection of BOLD responses and close agreement of the physiological components identified with Multicomponent PFM with those extracted with RETROICOR, even at low temporal resolution (TR 2s). As this technique is completely free of a-priori information it is ideally suited to the analysis of spontaneous brain activations in the resting state and experimental conditions where the timing of the BOLD signal changes cannot be predicted or measured.

This chapter finishes the presentation of the PFM techniques developed in this thesis. The next chapter will draw conclusions and outline future technical developments.

Chapter 7

Conclusions and further work

7.1 Summary and final discussion

One of the benefits of the increased signal to noise ratio (SNR) and contrast to noise ratio (CNR) available at ultra-high MRI systems (7 Tesla) is the possibility to study the brain's response to single trial events with BOLD fMRI, in contrast to the standard practice in event-related fMRI where the haemodynamic response is averaged across repeated events or stimuli. Single trial fMRI analysis has been typically performed based on prior knowledge about the timing of the events [54, 249, 298, 300]. Under certain clinical, experimental or behavioural conditions, it might be difficult to hypothesize a temporal model of cortical activations. This thesis has introduced novel fMRI analysis methods to study true single trial BOLD responses without prior information of when the event occurred: Paradigm Free Mapping.

To begin with, the principles of nuclear magnetic resonance, magnetic resonance imaging and functional magnetic resonance imaging were discussed in Chapters 2 and 3. The sensitivity to detect single trial BOLD responses in gradient echo (GE) fMRI experiments can significantly vary between brain regions owing to variations in T_2^* due to B_0 inhomogeneity and signal drop out due to B_1 inhomogeneity, particularly at high magnetic fields. BOLD fMRI sensitivity depends on both the signal change due to the BOLD effect and the temporal signal-to-noise ratio. In Chapter 3 a simple and novel method was developed to generate sensitivity maps showing, on a voxel-by-voxel basis, the change in T_2^* required to detect statistically significant activations. In addition, statistical correction factors were formulated to compensate for differences in T_2^* , which prevent optimal and uniform detection of the BOLD response, across brain regions. These sensitivity maps and correction factors provide

additional information for interpreting fMRI statistical parametric maps.

The motivations and principles behind Paradigm Free Mapping (PFM) were discussed in Chapter 4. Existing fMRI analysis methods to perform single trial experiments and which avoid precise or null specification of the experimental paradigm were also discussed in order to define a framework for the PFM techniques developed in this thesis: PFM (Chapter 4), Sparse PFM (Chapter 5) and Multicomponent PFM (Chapter 6). We aimed to merge both distinctive features: detect single trial BOLD responses by means of paradigm free analysis.

PFM methods are based on the deconvolution of the neuronal-related activity that drives the BOLD response assuming a haemodynamic response model. The L_2 -norm regularized estimator of Ridge Regression was studied for PFM in Chapter 4, whereas the LASSO (L_1 -norm regularized) and the Dantzig Selector (L_∞ -norm regularized) were investigated for Sparse PFM in Chapter 5. However, the performance of PFM techniques can deteriorate due to large baseline fluctuations of the signal. In Chapter 6, we presented a multicomponent PFM technique which decomposes the fMRI signal into a haemodynamic component describing the BOLD responses and a baseline component that can be described with sinusoidal basis functions. This multicomponent model was fitted using the morphological component analysis (MCA) algorithm, which proposes recursive L_1 -norm regularized estimation of both components of the fMRI signal.

Since the deconvolution is done by means of regularized or penalized estimators, a good selection of the regularization parameter λ is decisive to achieve a satisfactory operation of the PFM techniques. The selection of λ was based on Bayesian arguments for PFM, whereas model selection criteria (Akaike and Bayesian information criteria) and universal thresholds were recommended for Sparse PFM and Multicomponent PFM.

In terms of methodology, serial correlations of the noise were only considered for PFM in Chapter 4. The noise serial correlations were modelled by means of a parametric stochastic autoregressive model whose parameters were estimated with the Levinson Durbin algorithm. The assumption of uncorrelated noise in the models for Sparse and Multicomponent PFM was made to simplify computation. There remains the option to pre-whiten the data prior to analysis or incorporate a parametric stochastic model into the formulation.

Furthermore, the PFM techniques also differed in the way statistical inference was performed. Since ridge regression does not shrink the coefficient estimates to

zero, the significance of the deconvolved signal was assessed by means of a t -statistic against a baseline period in PFM (Chapter 4). In contrast, sparse estimation techniques effectively select the most relevant subset of coefficients describing the BOLD responses and shrink the rest of coefficients to zero. This remarkable feature aids statistical analysis since sparse estimation methods may ultimately permit that posterior thresholding of the model estimates is no longer necessary.

Simulated and experimental fMRI data were used to evaluate the performance of the techniques. We aimed to generate realistic simulated fMRI data to examine Sparse and Multicomponent PFM. The results of the simulations demonstrated that setting λ according to the Bayesian Information Criterion (BIC) or the Universal Threshold (UT) gave the best performance in Sparse PFM in order to obtain control over the false positive rate, whereas a choice based on the UT is advisable for Multicomponent PFM. It has been demonstrated that if λ is appropriately chosen, Sparse PFM and Multicomponent PFM can exhibit higher specificity values to detect the simulated events in time, and higher sensitivity when the model of the HRF is accurately known. However, the temporal sensitivity provided by both techniques diminished with a mismatch between the modelled and actual HRF. An important result of the simulations was the robustness of the technique against HRF-model mismatches in terms of the mean square error of the estimation of the BOLD responses existing in the fMRI signal despite the loss in temporal sensitivity of the techniques.

A visuomotor paradigm where the subject had to perform visually-cued and self-paced single trial finger tapping events served as a benchmark against which the three methods were experimentally evaluated. Electromyography (EMG) signals were simultaneously recorded to capture the timing of the events, and the three PFM methods were able to detect both visually-cued and self-paced events independently from the EMG signals. PFM showed higher resemblance in interpretation and results to standard GLM techniques than to ICA, suggesting that the PFM methodology is a valid solution to extend GLM-based approaches into a paradigm free formulation. Furthermore, Multicomponent PFM allowed us to identify low-frequency cardiac and respiratory fluctuations of the fMRI signal, showing similar performance to RETROICOR without the need of recording extra physiological signals.

In Chapter 5, the potential use of Sparse PFM to detect interictal epileptiform discharges (IED) was examined in datasets acquired in a patient with idiopathic generalized epilepsy. It must be emphasized that the goal of this study was not to

claim that SPFM can detect IEDs without the need to record EEG data, but that PFM techniques can assist to detect additional epileptogenic spikes and waves which might not be observed in the scalp EEG data. This could help to reduce the intra-subject and inter-subject variability that is commonly observed in EEG-based fMRI studies of epilepsy. Taking this into account, SPFM allowed us to detect 7 out of 8 of the EEG-marked IEDs in a paradigm free basis. Yet, additional activation events were also found with Sparse PFM showing signal changes in subcortical structures that are typically involve in interictal epileptic seizures, such as the thalamus and the caudate nucleus, calling for further investigation about the physiological origin of these activations in combination with other functional imaging techniques.

7.2 Further developments

The research presented in this thesis raises several points that deserve further investigation. To begin with, this thesis has demonstrated the great potential of using sparse estimators to deconvolve the BOLD responses without prior information of their timing. The first goal when using regularized estimators is to make an optimal selection of the regularization parameter, λ . In addition to the approaches investigated in this thesis, we will initially study alternative procedures where the regularization parameter is selected according to regional or tissue noise characteristics.

The next major difficulty in the application of Sparse PFM is to perform statistical inference of the estimates. This task requires studying the asymptotic properties of the sparse estimators [116, 383]. We proposed for Sparse PFM to approximate statistical inference based on t -statistics of the nonzero estimated coefficients. Nevertheless, defining the hypothesis of the activations based on the estimation results is certainly suboptimal and it is only valid for exploratory purposes. The validity and diagnosis of the estimated model can be done with diagnostic summary statistics based on the model residuals [232]. Alternatively, reformulating the ideas presented in this thesis into a Bayesian perspective [136, 131] might also be helpful to provide statistical inference of the activations. Estimators based on sparse Bayesian learning theories [367] have been shown equivalent or superior to L_1 -norm regularized estimators [366].

Methodologically, it will be also advisable to extend the methods into a hierarchical model so that the PFM algorithms can accommodate multiple data levels,

e.g. from multiple runs, sessions or subjects [129]. In addition, developing a spatio-temporal PFM approach will certainly improve the quality of the estimates and enhance the detection of the BOLD responses [354, 121, 281].

The major difficulty when performing a PFM analysis is to assess the functional relevance of the activations detected in the resting state. Although reverse mapping can be used to hypothesize that a certain pattern of activations is triggered by a cognitive task based on existing functional imaging studies, it will certainly be interesting to acquire EEG and fMRI data simultaneously. In addition to complement the high spatial resolution achievable with fMRI with the high temporal resolution provided by EEG, we can exploit EEG data to investigate which features of the EEG signal, if any, trigger the single trial haemodynamic responses detected with PFM [18, 93, 273]. This will ultimately help us to reduce the number of false positive activation events detected with PFM, particularly when the PFM activations occur in cortical areas near the scalp which can reliably be mapped with EEG source localization algorithms [252].

In terms of experimental paradigms, we are especially interested in applying PFM in the study of the dynamics of functional connectivity in the resting state [73] or cognitive experiments with spontaneous signal changes and where the timing of the brain activations cannot be established in advance, such as visual priming [356] or binocular rivalry [341]. In addition, our good preliminary results for the detection and characterization of interictal epileptiform discharges and epileptiform networks without the use of EEG data allow us to aspire to apply PFM techniques in epilepsy studies. From an MRI-engineering point of view, it will be stimulating to discover the highest spatial resolution at which event-related single-trial BOLD responses can be detected without using prior information about the event timing.

Finally, all the research and algorithms presented in this thesis have been conducted using code written in Matlab (MathWorks, Natick, MA) for personal use. To spread the application of the PFM techniques developed in this thesis, it will be useful to make the code available to other researchers as an extension or plugin to one of the widespread fMRI software packages, such as SPM (FIL, UCL, UK) [129], FSL (FMRIB, Oxford, UK) [324] or AFNI (NIMH/NIH, Bethesda, MD, USA) [82]. In any case, running Sparse PFM and Multicomponent PFM is computationally demanding. Therefore, translating the code of the proposed algorithms into a faster programming language, such as C++, will be one of the first tasks to undertake in the near future.

Finally, the Multicomponent PFM technique described in this thesis has demonstrated that the high CNR available at ultrahigh MRI allows the decomposition of the fMRI signal into multiple sources based on morphological characteristics. I believe that the fMRI community has, or will have, sufficient information about the origin, location, pattern and shape of the different components comprising the observed fMRI signal. Most of the information available about the physiological origin of the BOLD signal, noise and artefactual components is nowadays neglected in standard data-driven approaches. Therefore, it will be interesting to pursue a visionary decomposition of the fMRI signal by means of a spatio-temporal multicomponent, physiologically based, paradigm free mapping technique.

7.3 List of publications

Part of the work and results presented in this thesis have been presented in peer-review journal publications and international meetings:

- Caballero-Gaudes, C., Petridou, N., Dryden, I., Bai, L., Francis, S., and Gowland, P., Detection and Characterization of single-trial fMRI BOLD responses: Paradigm-Free Mapping. Accepted for publication *Human Brain Mapping*, May 2010.
- Caballero-Gaudes, C., Petridou, N., Francis, S., and Gowland, P., Unsupervised semiparametric Paradigm Free Mapping of BOLD responses and physiological fluctuations. *17th Annual Meeting of the Organization for Human Brain Mapping, HBM 2010*, Barcelona, Spain.
- Caballero-Gaudes, C., Petridou, N., Francis, S., and Gowland, P., Semiparametric Paradigm Free Mapping: Automatic detection and characterization of fMRI BOLD responses and physiological fluctuations without prior information. *18th International Society of Magnetic Resonance in Medicine, ISMRM 2010*, Stockholm, Sweden.
- Caballero-Gaudes, C., Petridou, N., Francis, S., Dryden, I., and Gowland, P., Sparse estimation using the Dantzig Selector algorithm automatically detects single-trial BOLD responses in space and time. *16th Annual Meeting of the Organization for Human Brain Mapping, HBM 2009*, San Francisco, California, USA.

- Caballero-Gaudes, C., Petridou, N., Francis, S., Driver, I., Dryden, I., and Gowland, P., Assessing BOLD sensitivity: implications at high field. *16th Annual Meeting of the Organization for Human Brain Mapping, HBM 2009*, San Francisco, California, USA.
- Petridou, N., Caballero-Gaudes, C., Dryden, I., Francis, S., and Gowland, P., Spatiotemporally coordinated activation detected during apparent rest in fMRI. *16th Annual Meeting of the Organization for Human Brain Mapping, HBM 2009*, San Francisco, California, USA.
- Petridou, N., Caballero-Gaudes, C., Dryden, I., Francis, S., and Gowland, P., Spatiotemporally coordinated activation detected during apparent rest in fMRI. *17th Annual Meeting of the International Society of Magnetic Resonance in Medicine, ISMRM 2009*, Honolulu, Hawaii, USA.
- Caballero-Gaudes, C., Petridou, N., Francis, S., Dryden, I., Bai, L., and Gowland, P., Detection of single- trial events in BOLD fMRI without prior stimulus information. *16th Annual Meeting of the International Society of Magnetic Resonance in Medicine, ISMRM 2008*, Toronto, Canada.
- Caballero-Gaudes, C., Petridou, N., Francis, S., Dryden, I., Bai, L., and Gowland, P., Detection of single- trial events in BOLD fMRI without prior stimulus information. *15th Annual Meeting of the Organization for Human Brain Mapping, HBM 2008*, Melbourne, Australia.

Appendices

A.1 Wiener Deconvolution

The ridge regression (RR) estimate can be interpreted as the maximum a-posteriori estimate (MAP) of a Bayesian estimator, i.e. the mode of the posterior distribution, of the coefficients \mathbf{s} given the data and the model parameters. Being the likelihood function $p(\mathbf{y}|\mathbf{s}, \sigma, \mathbf{\Sigma})$

$$p(\mathbf{y}|\mathbf{s}, \mathbf{\Sigma}) = (2\pi\sigma^2)^{-N/2} |\mathbf{\Sigma}|^{-1/2} \exp\left(-\frac{1}{2\sigma^2} (\mathbf{y} - \mathbf{H}\mathbf{s})^T \mathbf{\Sigma}^{-1} (\mathbf{y} - \mathbf{H}\mathbf{s})\right)$$

and setting a zero-mean uncorrelated Gaussian prior for \mathbf{s} , i.e. $\mathbf{s} \sim N(0, \beta^2\mathbf{I})$,

$$p(\mathbf{s}) = (2\pi\beta^2)^{-N/2} \exp\left(-\frac{\mathbf{s}^T \mathbf{s}}{2\beta^2}\right), \quad (1)$$

it can be shown that the posterior distribution of \mathbf{s} is $N(\mu, \mathbf{\Sigma}_s)$, where the mean is [41]

$$\mu = \left(\mathbf{H}^T \mathbf{\Sigma}^{-1} \mathbf{H} + \frac{\sigma^2}{\beta^2} \mathbf{I}\right)^{-1} \mathbf{H}^T \mathbf{\Sigma}^{-1} \mathbf{y}, \quad (2)$$

and the correlation matrix of the estimates is [41]

$$\mathbf{\Sigma}_s = \sigma^2 \left(\mathbf{H}^T \mathbf{\Sigma}^{-1} \mathbf{H} + \frac{\sigma^2}{\beta^2} \mathbf{I}\right)^{-1}. \quad (3)$$

Since the posterior distribution is Gaussian, the mode of the distribution is equal to its mean μ . This solution is also known as Wiener estimate or Wiener deconvolution [41, 142, 143]. Observe that the RR estimate

$$\mathbf{s}_{RR} = (\mathbf{H}^T \mathbf{\Sigma}^{-1} \mathbf{H} + \lambda \mathbf{I})^{-1} \mathbf{H}^T \mathbf{\Sigma}^{-1} \mathbf{y},$$

is equal to the Wiener estimate when $\lambda = \sigma^2/\beta^2$.

A.2 Prewhitening

This appendix demonstrates how the traditional concept of prewhitening is inherently included in the generalized least squares and ridge regression estimators.

Typically, prewhitening the data is understood as the temporal filtering of the data via matrix multiplication with the matrix \mathbf{S} [55, 369, 373]

$$\begin{aligned}\mathbf{S}\mathbf{y} &= \mathbf{S}\mathbf{H}\mathbf{s} + \mathbf{S}\boldsymbol{\epsilon} \\ \tilde{\mathbf{y}} &= \tilde{\mathbf{H}}\mathbf{s} + \tilde{\boldsymbol{\epsilon}},\end{aligned}\tag{4}$$

where $\tilde{\mathbf{y}} = \mathbf{S}\mathbf{y}$ and equivalently for $\tilde{\mathbf{H}}$ and $\tilde{\boldsymbol{\epsilon}}$. The autocorrelation matrix of $\tilde{\boldsymbol{\epsilon}}$ is $\mathbf{S}\boldsymbol{\Sigma}\mathbf{S}^T$, where $\boldsymbol{\Sigma}$ is the autocorrelation matrix of $\boldsymbol{\epsilon}$. Being $\boldsymbol{\Sigma}$ an autocorrelation matrix, there exists a square, nonsingular matrix such that $\boldsymbol{\Sigma} = \mathbf{K}\mathbf{K}^T$. By setting $\mathbf{S} = \mathbf{K}^{-1}$, one can easily see that the noise becomes uncorrelated, i.e. $\mathbf{K}^{-1}\boldsymbol{\Sigma}\mathbf{K}^{-T} = \mathbf{I}$. Hence, the OLS estimator in (4.10) can be used to compute the best linear unbiased estimates of s as

$$\hat{\mathbf{s}}_{OLS} = \left(\tilde{\mathbf{H}}^T\tilde{\mathbf{H}}\right)^{-1}\tilde{\mathbf{H}}^T\tilde{\mathbf{y}}.\tag{5}$$

Substituting $\tilde{\mathbf{H}} = \mathbf{S}\mathbf{H} = \mathbf{K}^{-1}\mathbf{H}$ and equivalently for $\tilde{\mathbf{y}}$ yields to the expression for the GLS estimator (4.9)

$$\hat{\mathbf{s}}_{GLS} = \left(\mathbf{H}^T\boldsymbol{\Sigma}^{-1}\mathbf{H}\right)^{-1}\mathbf{H}^T\boldsymbol{\Sigma}^{-1}\mathbf{y}.\tag{6}$$

and the RR estimator (4.14) by extending with the l_2 -norm regularization.

A.3 Levinson Durbin algorithm

The basic idea of the Levinson Durbin (LD) algorithm is to recursively solve the Yule-Walker equations, which define the autocorrelation function of an AR process, and estimate the AR parameters for each candidate order i ($a_i, i = 1, \dots, p$) starting from $i=1$. In the following, the Yule-Walker equations and the LD algorithm are described based on [333].

Yule-Walker equations

We recall that

$$\epsilon(n) + a_1\epsilon(n-1) + a_2\epsilon(n-2) + \dots + a_p\epsilon(n-p) = w(n) \quad (7)$$

define an AR model of order p . Multiplying by $\epsilon(n-k)$ and taking expectations, the autocorrelation function for this AR model is given by

$$r(k) + \sum_{i=1}^p a_i r(k-i) = \begin{cases} \sigma^2 & \text{for } k=0 \\ 0 & \text{for } k \neq 0 \end{cases} \quad (8)$$

This system of equations can be rewritten for a given order i as

$$\begin{bmatrix} r(0) & r(1) & \dots & r(i) \\ r(1) & r(0) & & \vdots \\ \vdots & & \ddots & r(1) \\ r(i) & \dots & & r(0) \end{bmatrix} \begin{bmatrix} 1 \\ a_1 \\ \vdots \\ a_i \end{bmatrix} = \begin{bmatrix} \sigma^2 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (9)$$

These equations are known as the Yule-Walker equations. In order to highlight the dependence of the Yule-Walker equations on the order i , we can rewrite (9) as

$$\mathbf{R}_{i+1} \begin{bmatrix} 1 \\ \theta_i \end{bmatrix} = \begin{bmatrix} \sigma_i^2 \\ \mathbf{0} \end{bmatrix} \quad (10)$$

with obvious definition for \mathbf{R}_{i+1} , and defining the vector $\theta_i = [a_1, \dots, a_i]^T$. Using only those equations in (9) equal to 0, it can be written

$$\begin{bmatrix} r(1) \\ \vdots \\ r(i) \end{bmatrix} + \begin{bmatrix} r(0) & r(1) & \dots & r(i-1) \\ \vdots & & \ddots & \vdots \\ r(i-1) & \dots & & r(0) \end{bmatrix} \begin{bmatrix} a_1 \\ \vdots \\ a_i \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}, \quad (11)$$

or simply $\mathbf{r}_i + \mathbf{R}_i\theta = \mathbf{0}$ such that the AR model parameters are given by

$$\theta = \mathbf{R}_i^{-1}\mathbf{r}_i. \quad (12)$$

Finally, the residual variance of the model is computed from (8) for $k=0$.

Levinson-Durbin algorithm

The Yule-Walker equations (12) need to be solved for each candidate order AR(i) $i = 0, \dots, p$, where the simplest model AR(0) corresponds to the case of uncorrelated noise and has not AR model coefficients. The main idea of the Levinson-Durbin algorithm is to solve the Yule-Walker equations recursively starting from the solution for $p=1$. From (10), we could update the Yule-Walker equations for the next order $i + 1$ as follows

$$\mathbf{R}_{i+2} \begin{bmatrix} 1 \\ \theta_i \\ 0 \end{bmatrix} = \left[\begin{array}{c|c} \mathbf{R}_{i+1} & r(i+1) \\ \hline r(i+1) & \tilde{\mathbf{r}}_i^T \\ \hline \tilde{\mathbf{r}}_i^T & r(0) \end{array} \right] \begin{bmatrix} 1 \\ \theta_i \\ 0 \end{bmatrix} = \begin{bmatrix} \sigma_i^2 \\ \mathbf{0} \\ \alpha_i \end{bmatrix}, \quad (13)$$

where $\tilde{\mathbf{r}}_i$ is defined as the reversal of a vector \mathbf{r}_i as follows,

$$\tilde{\mathbf{r}}_i = [r(i), \dots, r(1)]^T \quad (14)$$

and

$$\alpha_i = r(i+1) + \tilde{\mathbf{r}}_i^T \theta_i. \quad (15)$$

Note that (13) would be equal to (10) if $\alpha_i=0$ except for order $i + 1$. In order to force that, we define

$$k_{i+1} = -\frac{\alpha_i}{\sigma_i^2} = \frac{r(i+1) + \tilde{\mathbf{r}}_i^T \theta_i}{\sigma_i^2}. \quad (16)$$

From (13), we can develop the following system of equations

$$\mathbf{R}_{i+2} \left\{ \begin{bmatrix} 1 \\ \theta_i \\ 0 \end{bmatrix} + k_{i+1} \begin{bmatrix} 0 \\ \tilde{\theta}_i \\ 1 \end{bmatrix} \right\} = \begin{bmatrix} \sigma_i^2 \\ 0 \\ \alpha_i \end{bmatrix} + k_{i+1} \begin{bmatrix} \alpha_i \\ 0 \\ \sigma_i^2 \end{bmatrix} = \begin{bmatrix} \sigma_i^2 + k_{i+1} \alpha_i \\ 0 \\ 0 \end{bmatrix} \quad (17)$$

where we have used (11), (15) and (16).

Nevertheless, since the Yule-Walker equations (13) are unique for a given order i , (17) must correspond to the system of equations of the next order $i + 1$ such that

$$\mathbf{R}_{i+2} \begin{bmatrix} 1 \\ \theta_{i+1} \end{bmatrix} = \begin{bmatrix} \sigma_{i+1}^2 \\ 0 \end{bmatrix}. \quad (18)$$

Therefore, the AR model parameters of order $i + 1$ can be recursively computed

as

$$\theta_{i+1} = \begin{bmatrix} \theta_i \\ 0 \end{bmatrix} + k_{i+1} \begin{bmatrix} \tilde{\theta}_i \\ 1 \end{bmatrix}, \quad (19)$$

and the residual variance of $w(n)$ is

$$\sigma_{i+1}^2 = \sigma_i^2 (1 - |k_{i+1}|^2). \quad (20)$$

The reflection coefficients k_i are usually known as the reflection coefficients of the LD algorithm. In addition, $-k_i$, define the partial correlation coefficients which measure the correlation between the observations $\epsilon(n)$ and $\epsilon(n-i)$ after the correlation due to the intermediate values $\epsilon(n-1)$, $\epsilon(n-2)$, \dots , $\epsilon(n-i+1)$ has been eliminated [333]. The partial correlation coefficients can be used to estimate the optimal order of the model, \hat{p} , such that $k_{\hat{p}}$ is not significantly different from zero [369] and this approach is very similar to the AIC model selection criteria [49].

A.4 Autocorrelation coefficients from the AR and reflection coefficients

The autocorrelation coefficients which define Σ can be calculated from the AR coefficients and the reflection coefficients obtained during the Levinson-Durbin algorithm. Initializing with $r(0) = 1$ and $r(1) = -k_1$, the subsequent autocorrelation coefficients are computed recursively as

$$\begin{aligned} r_{i+1} &= -\tilde{\mathbf{r}}_i^T \left[\theta_i + k_{i+1} \tilde{\theta}_i \right] - k_{i+1}, & i = 1, \dots, p-1 \\ r_i &= -a_1 r(i-1) - \dots - a_p r(i-p), & i = p, \dots, N-1 \end{aligned} \quad (21)$$

A.5 Number of degrees of freedom of the t -statistic

Let $s_k(i)$ be the observation (ridge regression estimate) at voxel k at time i , and write

$$s_L^k(i) = \frac{1}{L} \sum_{i \in N(k)} s_l(i), \quad (22)$$

for the arithmetic mean at time i over the neighbourhood $N(k)$ of voxel k which consists of L voxels ($i = 1, \dots, N$; $k = 1, \dots, L$). Hereinafter we do not consider the superscript index k to simplify the notation. Consider the model where s_L are

jointly multivariate Normal with distributions during the baseline period given by

$$\hat{s}_L(i) \sim N(\mu_0, \sigma_L^2), \quad 1 \leq i \leq B \quad (23)$$

and the correlation matrix during the baseline period is the $B \times B$ matrix Σ_L . In addition, after the baseline period the marginal distribution at time i is

$$\hat{s}_L(i) \sim N(\mu_1, \sigma_L^2), \quad i > B \quad (24)$$

which is independent of the values in the baseline period.

We wish to test

$$H_0 : \mu_0 = \mu_1 \quad \text{versus} \quad H_1 : \mu_0 \neq \mu_1. \quad (25)$$

Let

$$\hat{\mu}_L = \frac{1}{B} \sum_{i=1}^B \hat{s}_L(i) \quad (26)$$

and

$$\hat{\sigma}_L^2 = \frac{1}{B-1} (\hat{\mathbf{s}}_L^B - \hat{\mu}_L \mathbf{1}_B)^T \Sigma_L^{-1} (\hat{\mathbf{s}}_L^B - \hat{\mu}_L \mathbf{1}_B) \quad (27)$$

be the baseline mean and variance estimates respectively, where $\hat{\mathbf{s}}_L^B$ is the B -vector of the spatially-averaged deconvolved time series for the baseline and $\mathbf{1}_B$ is a B -vector of ones. The correlation matrix Σ_L is estimated by standardizing the spatial average of the covariance matrices of the RR estimates for each baseline, i.e. we estimate the correlation matrix by standardizing (to unit diagonal elements)

$$\frac{1}{L} \sum_{l \in N(k)} (\mathbf{H}^T \Sigma^{-1} \mathbf{H} + \lambda_l \mathbf{I})^{-1}, \quad (28)$$

and λ_l is the regularization parameter for voxel l .

Under H_0 , we have for $i > B$,

$$\left(1 + \frac{1}{B}\right)^{-1/2} (\hat{s}_L(i) - \hat{\mu}_L) \sim N(0, \sigma_L^2), \quad (29)$$

and independently

$$\hat{\sigma}_L^2 \sim \frac{\sigma_L^2}{B-1} \chi_{B-1}^2, \quad (30)$$

applying Cochran's theorem [171] and treating Σ_L as known. Hence, it can be

written that

$$t(i) = \frac{\left(1 + \frac{1}{B}\right)^{-1/2} (\hat{s}_L(i) - \hat{\mu}_L)}{\hat{\sigma}_L} \sim t_{B-1}, \quad (31)$$

if H_0 is true, where t_ν is the Student's t -distribution with ν degrees of freedom. Here, we used the definition of the t -statistic such that if two random variables X and Y are independently distributed as $X \sim N(0, 1)$ and $Y \sim \chi_N^2$, then $X\sqrt{N}/\sqrt{Y} \sim t_N$.

Bibliography

- [1] A. Abou-Elseoud, T. Starck, J. Remes, J. Nikkinen, O. Tervonen, and V. Kiviniemi. The effect of model order selection in group PICA. *Human Brain Mapping*, 31(8):1207–1216, Aug 2010.
- [2] Y. Aghakhani, A. P. Bagshaw, C. G. Bénar, C. Hawco, F. Andermann, F. Dubeau, and J. Gotman. fMRI activation during spike and wave discharges in idiopathic generalized epilepsy. *Brain*, 127(5):1127–1144, May 2004.
- [3] Y. Aghakhani, E. Kobayashi, A. P. Bagshaw, C. Hawco, C. G. Bénar, F. Dubeau, and J. Gotman. Cortical and thalamic fMRI responses in partial epilepsy with focal and bilateral synchronous spikes. *Clinical Neurophysiology*, 117(1):177–191, Jan 2006.
- [4] G. K. Aguirre, E. Zarahn, and M. D’esposito. The variability of human, BOLD hemodynamic responses. *Neuroimage*, 8(4):360–369, Nov 1998.
- [5] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19:716–723, 1974.
- [6] P. J. Allen, O. Josephs, and R. Turner. A method for removing imaging artifact from continuous EEG recorded during functional MRI. *Neuroimage*, 12(2):230–239, Aug 2000.
- [7] E. Amaldi and V. Kann. On the approximability of minimizing nonzero variables or unsatisfied relations in linear systems. *Theoretical Computer Science*, 209(1-2):237–260, 1998.
- [8] A. H. Andersen, D. M. Gash, and M. J. Avison. Principal component analysis of the dynamic response measured by fMRI: A generalized linear systems framework. *Magnetic Resonance Imaging*, 17(6):795–815, Jul 1999.
- [9] A. Andrade, A. L. Paradis, S. Rouquette, and J. B. Poline. Ambiguous results in functional neuroimaging data analysis due to covariate correlation. *Neuroimage*, 10:483–486, 1999.
- [10] A. Antoniadis and J. Fan. Regularization of wavelet approximations. *Journal of the American Statistical Association*, 96(455):939–955, 2001.
- [11] J. S. Archer, D. F. Abbott, A. B. Waites, and G. D. Jackson. fMRI “deactivation” of the posterior cingulate during generalized spike and wave. *Neuroimage*, 20(4):1915–1922, Dec 2003.
- [12] J. S. Archer, A. B. Waites, D. F. Abbott, P. Federico, and G. D. Jackson. Event-related fMRI of myoclonic jerks arising from dysplastic cortex. *Epilepsia*, 47(9):1487–1492, Sep 2006.
- [13] K. Arfanakis, D. Cordes, V. M. Haughton, C. H. Moritz, M. A. Quigley, and M. E. Meyerand. Combining Independent Component Analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. *Magnetic Resonance Imaging*, 18(8):921–930, Oct 2000.

- [14] A. Arieli, D. Shoham, R. Hildesheim, and A. Grinvald. Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *Journal of Neurophysiology*, 73(5):2072–2093, May 1995.
- [15] M. Salman Asif and J. Romberg. On the LASSO and Dantzig Selector equivalence. In *Conference on Information Sciences and Systems (CISS)*, 2010.
- [16] T. Auer, K. Veto, T. Doczi, S. Komoly, V. Juhos, J. Janszky, and A. Schwarcz. Identifying seizure-onset zone and visualizing seizure spread by fMRI: A case report. *Epileptic Disorders*, 10(2):93–100, Jun 2008.
- [17] A. P. Bagshaw, Y. Aghakhani, C. G. Bénar, E. Kobayashi, C. Hawco, F. Dubeau, G. B. Pike, and J. Gotman. EEG-fMRI of focal epileptic spikes: Analysis with multiple haemodynamic functions and comparison with gadolinium-enhanced MR angiograms. *Human Brain Mapping*, 22(3):179–192, 2004.
- [18] A. P. Bagshaw and T. Warbrick. Single trial variability of EEG and fMRI responses to visual stimuli. *Neuroimage*, 38(2):280–292, Nov 2007.
- [19] S. Baillet, J. C. Mosher, and R. M. Leahy. Electromagnetic brain mapping. *IEEE Signal Processing Magazine*, 18(6):14–30, Nov 2001.
- [20] C. J. Bakker, C. N. de Graaf, and P. van Dijk. Derivation of quantitative information in NMR imaging: A phantom study. *Physics in Medicine and Biology*, 29(12):1511–1525, Dec 1984.
- [21] E. Balteau, C. Hutton, and N. Weiskopf. Improved shimming for fMRI specifically optimizing the local BOLD sensitivity. *Neuroimage*, 49(1):327–336, Jan 2010.
- [22] P. A. Bandettini, E. C. Wong, R. S. Hinks, R. S. Tikofsky, and J. S. Hyde. Time course EPI of human brain function during task activation. *Magnetic Resonance in Medicine*, 25(2):390–397, Jun 1992.
- [23] P. Baraldi, A. A. Manginelli, M. Maieron, D. Liberati, and C. A. Porro. An ARX model-based approach to trial by trial identification of fMRI-BOLD responses. *Neuroimage*, 37(1):189–201, Aug 2007.
- [24] R. Baumgartner, L. Ryner, W. Richter, R. Summers, M. Jarmasz, and R. Somorjai. Comparison of two exploratory data analysis methods for fMRI: Fuzzy Clustering vs. Principal Component Analysis. *Magnetic Resonance Imaging*, 18(1):89–94, Jan 2000.
- [25] R. Baumgartner, R. Somorjai, R. Summers, and W. Richter. Assessment of cluster homogeneity in fMRI data using Kendall’s coefficient of concordance. *Magnetic Resonance Imaging*, 17(10):1525–1532, Dec 1999.
- [26] R. Baumgartner, R. Somorjai, R. Summers, W. Richter, L. Ryner, and M. Jarmasz. Resampling as a cluster validation technique in fMRI. *Journal of Magnetic Resonance Imaging*, 11(2):228–231, Feb 2000.

- [27] C. F. Beckmann and S. M. Smith. Probabilistic Independent Component Analysis for functional magnetic resonance imaging. *IEEE Transactions Medical Imaging*, 23(2):137–152, Feb 2004.
- [28] Y. Behzadi and T. T. Liu. Caffeine reduces the initial dip in the visual BOLD response at 3 T. *Neuroimage*, 32(1):9–15, Aug 2006.
- [29] A. J. Bell and T. J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159, Nov 1995.
- [30] P. S. F. Bellgowan, Z. S. Saad, and P. A. Bandettini. Understanding neural system dynamics through task modulation and measurement of functional MRI amplitude, latency, and width. *Proceedings of the National Academy of Sciences USA*, 100(3):1415–1419, Feb 2003.
- [31] J. W. Belliveau, D. N. Kennedy, R. C. McKinstry, B. R. Buchbinder, R. M. Weisskoff, M. S. Cohen, J. M. Vevea, T. J. Brady, and B. R. Rosen. Functional mapping of the human visual cortex by magnetic resonance imaging. *Science*, 254(5032):716–719, Nov 1991.
- [32] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 57(1):289–300, Jan 1995.
- [33] Y. Benjamini, A. M. Krieger, and D. Yekutieli. Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3):491–507, Sep 2006.
- [34] M. A. Bernstein, K. F. King, and X. J. Zhou. *Handbook of MRI Pulse Sequences*. Academic Press, 2004.
- [35] M. Bianciardi, M. Fukunaga, P. van Gelderen, S. G. Horovitz, J. A. de Zwart, K. Shmueli, and J. H. Duyn. Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: A 7 T study. *Magnetic Resonance Imaging*, 27(8):1019–1029, Oct 2009.
- [36] P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of LASSO and Dantzig Selector. *The Annals of Statistics*, 37(4):1705–1732, Aug 2009.
- [37] U. Bingel, M. Quante, R. Knab, B. Bromm, C. Weiller, and C. Büchel. Subcortical structures involved in pain processing: Evidence from single-trial fMRI. *Pain*, 99(1-2):313–321, Sep 2002.
- [38] R. M. Birn, J. B. Diamond, M. A. Smith, and P. A. Bandettini. Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *Neuroimage*, 31(4):1536–1548, Jul 2006.
- [39] R. M. Birn, Z. S. Saad, and P. A. Bandettini. Spatial heterogeneity of the nonlinear dynamics in the fMRI BOLD response. *Neuroimage*, 14(4):817–826, Oct 2001.
- [40] R. M. Birn, M. A. Smith, T. B. Jones, and P. A. Bandettini. The respiration response function: The temporal dynamics of fMRI signal fluctuations related to changes in respiration. *Neuroimage*, 40(2):644–654, Apr 2008.

- [41] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [42] F. Bloch, W.W. Hansen, and M. Packard. Nuclear induction. *Physical Review*, 70(7-8):460–474, 1946.
- [43] J. Bobin, J. L. Starck, J. M. Fadili, Y. Moudden, and D. L. Donoho. Morphological component analysis: An adaptive thresholding strategy. *IEEE Transactions on Image Processing*, 16(11):2675–2680, Nov 2007.
- [44] J. Bodurka, F. Ye, N. Petridou, K. Murphy, and P. A. Bandettini. Mapping the MRI voxel volume in which thermal noise matches physiological noise—implications for fMRI. *Neuroimage*, 34(2):542–549, Jan 2007.
- [45] A. Bolstad, B. Van Veen, and R. Nowak. Space-time event sparse penalization for magneto-/electroencephalography. *Neuroimage*, 46(4):1066–1081, Jul 2009.
- [46] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel. *Time Series Analysis: Forecasting and Control (Wiley Series in Probability and Statistics)*. Wiley, 4 edition, 2008.
- [47] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [48] G. M. Boynton, S. A. Engel, G. H. Glover, and D. J. Heeger. Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, 16(13):4207–4221, Jul 1996.
- [49] P. M. T. Broersen. *Automatic Autocorrelation and Spectral Analysis*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [50] A. M. Bruckstein, D.L. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, Feb 2009.
- [51] R. L. Buckner. Event-related fMRI and the hemodynamic response. *Human Brain Mapping*, 6(5-6):373–377, 1998.
- [52] R. L. Buckner. The hemodynamic inverse problem: making inferences about neural activity from measured MRI signals. *Proceedings of the National Academy of Sciences USA*, 100(5):2177–2179, Mar 2003.
- [53] R. L. Buckner, J. R. Andrews-Hanna, and D. L. Schacter. The brain’s default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, 1124(1):1–38, Mar 2008.
- [54] R. L. Buckner, P. A. Bandettini, K. M. O’Craven, R. L. Savoy, S. E. Petersen, M. E. Raichle, and B. R. Rosen. Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. *Proceedings of the National Academy of Sciences USA*, 93(25):14878–14883, Dec 1996.

- [55] E. Bullmore, M. Brammer, S. C. Williams, S. Rabe-Hesketh, N. Janot, A. David, J. Mellers, R. Howard, and P. Sham. Statistical methods of estimation and inference for functional MR image analysis. *Magnetic Resonance in Medicine*, 35(2):261–277, Feb 1996.
- [56] E. Bullmore, C. Long, J. Suckling, J. Fadili, G. Calvert, F. Zelaya, T. A. Carpenter, and M. Brammer. Colored noise and computational inference in neurophysiological fMRI time series analysis: Resampling methods in time and wavelet domains. *Human Brain Mapping*, 12(2):61–78, Feb 2001.
- [57] K. Butts, S.J. Riederer, R.L. Ehman, R.M. Thompson, and C.R. Jack. Interleaved echo planar imaging on a standard MRI systems. *Magnetic Resonance in Medicine*, 31(1):67–72, Jan 1994.
- [58] R. B. Buxton. The elusive initial dip. *Neuroimage*, 13(6 Pt 1):953–958, Jun 2001.
- [59] R. B. Buxton. *Introduction to Functional Magnetic Resonance Imaging : Principles and Techniques*. Cambridge University Press, 2nd edition, 2009.
- [60] C. G. Bénar, D. W. Gross, Y. Wang, V. Petre, B. Pike, F. Dubeau, and J. Gotman. The BOLD response to interictal epileptiform discharges. *Neuroimage*, 17(3):1182–1192, Nov 2002.
- [61] C. G. Bénar, C. Grova, E. Kobayashi, A. P. Bagshaw, Y. Aghakhani, F. Dubeau, and J. Gotman. EEG-fMRI of epileptic spikes: Concordance with EEG source localization and intracranial EEG. *Neuroimage*, 30(4):1161–1170, May 2006.
- [62] C. Büchel, A. P. Holmes, G. Rees, and K. J. Friston. Characterizing stimulus-response functions using nonlinear regressors in parametric fMRI experiments. *Neuroimage*, 8(2):140–148, Aug 1998.
- [63] C. Büchel, J. Morris, R. J. Dolan, and K. J. Friston. Brain systems mediating aversive conditioning: An event-related fMRI study. *Neuron*, 20(5):947–957, May 1998.
- [64] C. Caballero-Gaudes, I. Santamaria, J. Via, E. M. Gomez, and T. S. Paules. Robust array beamforming with sidelobe control using support vector machines. *IEEE Transactions Signal Processing*, 55(2):574–584, Feb. 2007.
- [65] V. D. Calhoun and T. Adali. Unmixing fMRI with Independent Component Analysis. *IEEE Engineering in Medicine and Biology Magazine*, 25(2):79–90, Mar-Apr 2006.
- [66] V. D. Calhoun, T. Adali, G. D. Pearlson, and J. J. Pekar. Spatial and temporal Independent Component Analysis of functional MRI data containing a pair of task-related waveforms. *Human Brain Mapping*, 13(1):43–53, May 2001.
- [67] V. D. Calhoun, M. C. Stevens, G. D. Pearlson, and K. A. Kiehl. fMRI analysis with the general linear model: Removal of latency-induced amplitude bias by incorporation of hemodynamic derivative terms. *Neuroimage*, 22(1):252–257, May 2004.

- [68] E. Candes and T. Tao. The Dantzig Selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6):2313–2351, Dec 2007.
- [69] J. F. Cardoso and A. Souloumiac. Blind beamforming for non gaussian signals. *IEE Proceedings-F*, 140(6):362–370, Dec 1993.
- [70] R. Casanova, S. Ryali, J. Serences, L. Yang, R. Kraft, P. J. Laurienti, and J. A. Maldjian. The impact of temporal regularization on estimates of the BOLD hemodynamic response function: A comparative analysis. *Neuroimage*, 40(4):1606–1618, May 2008.
- [71] C. Chang, J. P. Cunningham, and G. H. Glover. Influence of heart rate on the BOLD signal: The cardiac response function. *Neuroimage*, 44(3):857–869, Feb 2009.
- [72] C. Chang and G. H. Glover. Effects of model-based physiological noise correction on default mode network anti-correlations and correlations. *Neuroimage*, 47(4):1448–1459, Oct 2009.
- [73] C. Chang and G. H. Glover. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage*, 50(1):81–98, Mar 2010.
- [74] R. Chartrand. Exact reconstructions of sparse signals via nonconvex minimization. *IEEE Signal Processing Letters*, 14(10):707–710, Oct 2007.
- [75] J. J. Chen and G. B. Pike. Origins of the BOLD post-stimulus undershoot. *Neuroimage*, 46(3):559–568, Jul 2009.
- [76] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, Aug 1998.
- [77] J. Chumbley, K. Worsley, G. Flandin, and K. Friston. Topological FDR for neuroimaging. *Neuroimage*, 49(4):3057–3064, Feb 2010.
- [78] P. Ciuciu, J. B. Poline, G. Marrelec, J. Idier, Ch. Pallier, and H. Benali. Unsupervised robust nonparametric estimation of the hemodynamic response function for any fMRI experiment. *IEEE Transactions on Medical Imaging*, 22(10):1235–1251, Oct 2003.
- [79] M. S. Cohen. Parametric analysis of fMRI data using linear systems methods. *Neuroimage*, 6(2):93–103, Aug 1997.
- [80] M. Corbetta and G. L. Shulman. Control of goal-directed and stimulus-driven attention in the brain. *Nature Review Neuroscience*, 3(3):201–215, Mar 2002.
- [81] N. Correa, T. Adali, and V. D. Calhoun. Performance of blind source separation algorithms for fMRI analysis using a group ICA method. *Magnetic Resonance Imaging*, 25(5):684–694, Jun 2007.
- [82] R. W. Cox. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29(3):162–173, Jun 1996.

- [83] R. Cunnington, C. Windischberger, L. Deecke, and E. Moser. The preparation and execution of self-initiated and externally-triggered movement: A study of event-related fMRI. *Neuroimage*, 15(2):373–385, Feb 2002.
- [84] R. Cunnington, C. Windischberger, L. Deecke, and E. Moser. The preparation and readiness for voluntary movement: A high-field event-related fMRI study of the Bereitschafts-BOLD response. *Neuroimage*, 20(1):404–412, Sep 2003.
- [85] R. Cunnington, C. Windischberger, S. Robinson, and E. Moser. The selection of intended actions and the observation of others’ actions: A time-resolved fMRI study. *Neuroimage*, 29(4):1294–1302, Feb 2006.
- [86] M. S. Dagli, J. E. Ingeholm, and J. V. Haxby. Localization of cardiac-induced signal change in fMRI. *Neuroimage*, 9(4):407–415, Apr 1999.
- [87] Hanna Damasio. *Human brain anatomy in computerized images*. Oxford University Press, New York, 1995.
- [88] J. S. Damoiseaux, S. A. R. B. Rombouts, F. Barkhof, P. Scheltens, C. J. Stam, S. M. Smith, and C. F. Beckmann. Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences USA*, 103(37):13848–13853, Sep 2006.
- [89] I. Daubechies, E. Roussos, S. Takerkart, M. Benharrosh, C. Golden, K. D’Ardenne, W. Richter, J. D. Cohen, and J. Haxby. Independent component analysis for brain fMRI does not select for independence. *Proceedings of the National Academy of Sciences USA*, 106(26):10415–10422, Jun 2009.
- [90] P. F. Van de Moortele, J. Pfeuffer, G. H. Glover, K. Uğurbil, and X. Hu. Respiration-induced B0 fluctuations and their spatial distribution in the human brain at 7 Tesla. *Magnetic Resonance in Medicine*, 47(5):888–895, May 2002.
- [91] J. A. de Zwart, P. van Gelderen, and J. H. Duyn. Receive coil arrays and parallel imaging for functional magnetic resonance imaging of the human brain. *Conf Proc IEEE Eng Med Biol Soc*, 1:17–20, 2006.
- [92] J. A. de Zwart, P. van Gelderen, J. M. Jansma, M. Fukunaga, M. Bianciardi, and J. H. Duyn. Hemodynamic nonlinearities affect BOLD fMRI response timing and amplitude. *Neuroimage*, 47(4):1649–1658, Oct 2009.
- [93] S. Debener, M. Ullsperger, M. Siegel, and A. K. Engel. Single-trial EEG-fMRI reveals the dynamics of cognitive function. *Trends in Cognitive Sciences*, 10(12):558–563, Dec 2006.
- [94] S. Debener, M. Ullsperger, M. Siegel, K. Fiehler, D. Y. von Cramon, and A. K. Engel. Trial-by-trial coupling of concurrent electroencephalogram and functional magnetic resonance imaging identifies the dynamics of performance monitoring. *Journal of Neuroscience*, 25(50):11730–11737, Dec 2005.
- [95] R. Deichmann, J. A. Gottfried, C. Hutton, and R. Turner. Optimized EPI for fMRI studies of the orbitofrontal cortex. *Neuroimage*, 19(2 Pt 1):430–441, Jun 2003.

- [96] R. Deichmann, O. Josephs, C. Hutton, D. R. Corfield, and R. Turner. Compensation of susceptibility-induced BOLD sensitivity losses in echo-planar fMRI imaging. *Neuroimage*, 15(1):120–135, Jan 2002.
- [97] T. Deneux and O. Faugeras. EEG-fMRI fusion of paradigm-free activity using Kalman filtering. *Neural Computation*, 22(4):906–948, Apr 2010.
- [98] H. Ding, G. Claeskens, and M. Jansen. Variable selection in partially linear wavelet models. *Statistical Modelling*, in press, 2010.
- [99] A. Donaire, N. Bargallo, C. Falcón, I. Maestro, M. Carreno, J. Setoain, J. Rumià, S. Fernández, L. Pintor, and T. Boget. Identifying the structures involved in seizure generation using sequential analysis of ictal-fMRI data. *Neuroimage*, 47(1):173–183, Aug 2009.
- [100] A. Donaire, C. Falcón, M. Carreno, N. Bargallo, J. Rumià, J. Setoain, I. Maestro, T. Boget, L. Pintor, R. Agudo, M. Falip, and S. Fernández. Sequential analysis of fMRI images: A new approach to study human epileptic networks. *Epilepsia*, 50(12):2526–2537, Dec 2009.
- [101] D. L. Donoho and I. M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455, Sep 1994.
- [102] I. Driver, N. Blockley, J. Fisher, S. Francis, and P. Gowland. The change in cerebrovascular reactivity between 3 T and 7 T measured using graded hypercapnia. *Neuroimage*, 51:274–279, Jan 2010.
- [103] J. R. Duann, T. P. Jung, W. J. Kuo, T. C. Yeh, S. Makeig, J. C. Hsieh, and T. J. Sejnowski. Single-trial variability in event-related BOLD signals. *Neuroimage*, 15(4):823–835, Apr 2002.
- [104] T. Q. Duong, D. S. Kim, K. Uğurbil, and S. G. Kim. Spatiotemporal dynamics of the BOLD fMRI signals: Toward mapping submillimeter cortical columns using the early negative response. *Magnetic Resonance in Medicine*, 44(2):231–242, Aug 2000.
- [105] W. A. Edelstein, G. H. Glover, C. J. Hardy, and R. W. Redington. The intrinsic signal-to-noise ratio in NMR imaging. *Magnetic Resonance in Medicine*, 3(4):604–618, Aug 1986.
- [106] W.A. Edelstein, J.M.S. Hutchison, G.Johnson, and T.Redpath. Spin warp NMR imaging and applications to human whole-body imaging. *Physics in Medicine and Biology*, 25(4):751–756, Jul 1980.
- [107] B. Efron. The estimation of prediction error: Covariance penalties and cross-validation. *Journal of the American Statistical Association*, 99(467):619–632, Sep 2004.
- [108] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32(2):407–451, Apr 2004.
- [109] T. Eichele, V. D. Calhoun, M. Moosmann, K. Specht, M. L. A. Jongsma, R. Q. Quiroga, H. Nordby, and K. Hugdahl. Unmixing concurrent EEG-fMRI with parallel Independent Component Analysis. *International Journal of Psychophysiology*, 67(3):222–234, Mar 2008.

- [110] D. Erdogmus, K. E. Hild, J. C. Principe, M. Lazaro, and I. Santamaria. Adaptive blind deconvolution of linear channels using Renyi's entropy with Parzen window estimation. *IEEE Transactions on Signal Processing*, 52(6):1489–1498, Jun 2004.
- [111] D. Erdogmus, K. E. Hild, Y. N. Rao, and J. C. Príncipe. Minimax mutual information approach for Independent Component Analysis. *Neural Computation*, 16(6):1235–1252, Jun 2004.
- [112] R.R. Ernst and W.A. Anderson. Application of Fourier transform spectroscopy to magnetic resonance. *Review of Scientific Instruments*, 37(1):93, Jan 1966.
- [113] T. Ernst and J. Hennig. Observation of a fast response in functional MR. *Magnetic Resonance in Medicine*, 32(1):146–149, Jul 1994.
- [114] F. Esposito, E. Formisano, E. Seifritz, R. Goebel, R. Morrone, G. Tedeschi, and F. Di Salle. Spatial Independent Component Analysis of functional MRI time-series: To what extent do results depend on the algorithm used? *Human Brain Mapping*, 16(3):146–157, Jul 2002.
- [115] J.M. Fadili, J. L. Starck, J. Bobin, and Y. Moudden. Image decomposition and separation using sparse representations: An overview. *Proceedings of the IEEE*, 98(6):983–994, Jun 2010.
- [116] M. J. Fadili and E. T. Bullmore. Penalized partially linear models using sparse representations with an application to fMRI time series. *IEEE Transactions on Signal Processing*, 53(9):3436–3448, Sep 2005.
- [117] M. J. Fadili, S. Ruan, D. Bloyet, and B. Mazoyer. A multistep unsupervised Fuzzy Clustering analysis of fMRI time series. *Human Brain Mapping*, 10(4):160–178, Aug 2000.
- [118] S. Faisan, L. Thoraval, J. P. Armspach, and F. Heitz. Hidden Markov multiple event sequence models: A paradigm for the spatio-temporal analysis of fMRI data. *Medical Image Analysis*, 11(1):1–20, Feb 2007.
- [119] O. Felician and P. Romaguère. Your body and mine: A neuropsychological perspective. *Clinical Neurophysiology*, 38(3):183–187, Jun 2008.
- [120] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE Journal of Selected Topics in Signal Processing*, 1(4):586–597, Dec 2007.
- [121] G. Flandin and W. D. Penny. Bayesian fMRI data analysis with sparse spatial basis function priors. *Neuroimage*, 34(3):1108–1125, Feb 2007.
- [122] E. Formisano, F. Esposito, F. Di Salle, and R. Goebel. Cortex-based Independent Component Analysis of fMRI time series. *Magnetic Resonance Imaging*, 22(10):1493–1504, Dec 2004.
- [123] E. Formisano, D. E. J. Linden, F. Di Salle, L. Trojano, F. Esposito, A. T. Sack, D. Grossi, F. E. Zanella, and R. Goebel. Tracking the mind's image in the brain I: Time-resolved fMRI during visuospatial mental imagery. *Neuron*, 35(1):185–194, Jul 2002.

- [124] M. D. Fox and M. E. Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Review Neuroscience*, 8(9):700–711, Sep 2007.
- [125] M. D. Fox, A. Z. Snyder, J. L. Vincent, and M. E. Raichle. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, 56(1):171–184, Oct 2007.
- [126] P. T. Fox, M. A. Mintun, M. E. Raichle, F. M. Miezin, J. M. Allman, and D. C. Van Essen. Mapping human visual cortex with positron emission tomography. *Nature*, 323(6091):806–809, 1986.
- [127] R. S. Frackowiak and K. J. Friston. Functional neuroanatomy of the human brain: Positron emission tomography—a new neuroanatomical technique. *Journal of Anatomy*, 184 (Pt 2):211–225, Apr 1994.
- [128] S. Francis, X. Lin, S. Aboushoushah, T. P. White, M. Phillips, R. Bowtell, and C. S. Constantinescu. fMRI analysis of active, passive and electrically stimulated ankle dorsiflexion. *Neuroimage*, 44(2):469–479, Jan 2009.
- [129] K. J. Friston, J. Ashburner, S. Kiebel, T. Nichols, and W. D. Penny. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Elsevier, 1st edition, 2006.
- [130] K. J. Friston, P. Fletcher, O. Josephs, A. Holmes, M. D. Rugg, and R. Turner. Event-related fMRI: Characterizing differential responses. *Neuroimage*, 7(1):30–40, Jan 1998.
- [131] K. J. Friston, D. E. Glaser, R. N. A. Henson, S. Kiebel, C. Phillips, and J. Ashburner. Classical and Bayesian inference in neuroimaging: Applications. *Neuroimage*, 16(2):484–512, Jun 2002.
- [132] K. J. Friston, A. P. Holmes, J. B. Poline, P. J. Grasby, S. C. Williams, R. S. Frackowiak, and R. Turner. Analysis of fMRI time-series revisited. *Neuroimage*, 2(1):45–53, Mar 1995.
- [133] K. J. Friston, P. Zeigler, and R. Turner. Analysis of functional MRI time-series. *Human Brain Mapping*, 1:153–171, Jan 1994.
- [134] K. J. Friston, O. Josephs, G. Rees, and R. Turner. Nonlinear event-related responses in fMRI. *Magnetic Resonance in Medicine*, 39(1):41–52, Jan 1998.
- [135] K. J. Friston, O. Josephs, E. Zarahn, A. P. Holmes, S. Rouquette, and J. B. Poline. To smooth or not to smooth? Bias and efficiency in fMRI time-series analysis. *Neuroimage*, 12(2):196–208, Aug 2000.
- [136] K. J. Friston, W. Penny, C. Phillips, S. Kiebel, G. Hinton, and J. Ashburner. Classical and Bayesian inference in neuroimaging: Theory. *Neuroimage*, 16(2):465–483, Jun 2002.
- [137] K. J. Friston, S. Williams, R. Howard, R. S. J. Frackowiak, and R. Turner. Movement-related effects in fMRI time-series. *Magnetic Resonance in Medicine*, 35(3):346–355, Mar 1996.

- [138] K. J. Friston, E. Zarahn, O. Josephs, R. N. Henson, and A. M. Dale. Stochastic designs in event-related fMRI. *Neuroimage*, 10(5):607–619, Nov 1999.
- [139] A. N. Garroway, P. K. Grannell, and P. Mansfield. Image formation in NMR by a selective irradiative process. *Journal of Physics C: Solid State Physics*, 7(24):L457–62, Dec 1974.
- [140] T. Gautama and M. M. Van Hulle. Optimal spatial regularisation of auto-correlation estimates in fMRI analysis. *Neuroimage*, 23(3):1203–1216, Nov 2004.
- [141] C. R. Genovese, N. A. Lazar, and T. Nichols. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, 15(4):870–878, Apr 2002.
- [142] D. R. Gitelman, W. D. Penny, J. Ashburner, and K. J. Friston. Modeling regional and psychophysiologic interactions in fMRI: The importance of hemodynamic deconvolution. *Neuroimage*, 19(1):200–207, May 2003.
- [143] G. H. Glover. Deconvolution of impulse response in event-related BOLD fMRI. *Neuroimage*, 9(4):416–429, Apr 1999.
- [144] G. H. Glover, T. Q. Li, and D. Ress. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine*, 44(1):162–167, Jul 2000.
- [145] D. Gonzalo, T. Shallice, and R. Dolan. Time-dependent changes in learning audiovisual associations: A single-trial fMRI study. *Neuroimage*, 11(3):243–255, Mar 2000.
- [146] A. M. Gordon, J. H. Lee, D. Flament, K. Uğurbil, and T. J. Ebner. Functional magnetic resonance imaging of motor, sensory, and posterior parietal cortical areas during performance of sequential typing movements. *Experimental Brain Research*, 121(2):153–166, Jul 1998.
- [147] J. Gotman, C. Grova, A. Bagshaw, E. Kobayashi, Y. Aghakhani, and F. Dubeau. Generalized epileptic discharges show thalamocortical activation and suspension of the default state of the brain. *Proceedings of the National Academy of Sciences USA*, 102(42):15236–15240, Oct 2005.
- [148] J. Gotman, E. Kobayashi, A. P. Bagshaw, C. G. Bénar, and F. Dubeau. Combining EEG and fMRI: A multimodal tool for epilepsy research. *Journal of Magnetic Resonance Imaging*, 23(6):906–920, Jun 2006.
- [149] C. Goutte, F. A. Nielsen, and L. K. Hansen. Modeling the haemodynamic response in fMRI using smooth FIR filters. *IEEE Transactions on Medical Imaging*, 19(12):1188–1201, Dec 2000.
- [150] C. Goutte, P. Toft, E. Rostrup, F. Nielsen, and L. K. Hansen. On clustering fMRI time series. *Neuroimage*, 9(3):298–310, Mar 1999.
- [151] P. A. Gowland and R. Bowtell. Theoretical optimization of multi-echo fMRI data acquisition. *Physics in Medicine and Biology*, 52(7):1801–1813, Apr 2007.

- [152] G. Gratton and M. Fabiani. The event-related optical signal: A new tool for studying brain function. *International Journal of Psychophysiology*, 42(2):109–121, Oct 2001.
- [153] G. Gratton, M. Fabiani, P. M. Corballis, D. C. Hood, M. R. Goodman-Wood, J. Hirsch, K. Kim, D. Friedman, and E. Gratton. Fast and localized event-related optical signals (EROS) in the human occipital cortex: Comparisons with the visual evoked potential and fMRI. *Neuroimage*, 6(3):168–180, Oct 1997.
- [154] K. Grill-Spector and R. Malach. fMR-Adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica (Amsterdam)*, 107(1-3):293–321, Apr 2001.
- [155] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine*, 47(6):1202–1210, Jun 2002.
- [156] A. Haase, D. Matthaei, W. Hänicke, and J. Frahm. Dynamic digital subtraction imaging using fast low-angle shot MR movie sequence. *Radiology*, 160(2):537–541, Aug 1986.
- [157] E.L. Hahn. Spin echoes. *Physical Review*, 80(4):580–594, Nov 1950.
- [158] U. Halsband, N. Ito, J. Tanji, and H. J. Freund. The role of premotor cortex and the supplementary motor area in the temporal control of movement in man. *Brain*, 116 (Pt 1):243–266, Feb 1993.
- [159] M. Hämäläinen, R. Hari, R. J. Ilmoniemi, J. Knuutila, and O. V. Lounasmaa. Magnetoencephalography- theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics*, 65(2):413–497, Apr 1993.
- [160] K. Hamandi, A. Salek Haddadi, A. Liston, H. Laufs, D. R. Fish, and L. Lemieux. fMRI temporal clustering analysis in patients with frequent interictal epileptiform discharges: Comparison with EEG-driven analysis. *Neuroimage*, 26(1):309–316, May 2005.
- [161] K. Hamandi, A. Salek-Haddadi, H. Laufs, A. Liston, K. Friston, D. R. Fish, J. S. Duncan, and L. Lemieux. EEG-fMRI of idiopathic and secondarily generalized epilepsies. *Neuroimage*, 31(4):1700–1710, Jul 2006.
- [162] D.A. Handwerker, J. M. Ollinger, and M. D’Esposito. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage*, 21(4):1639–1651, Apr 2004.
- [163] L. K. Hansen, J. Larsen, F. A. Nielsen, S. C. Strother, E. Rostrup, R. Savoy, N. Lange, J. Sidtis, C. Svarer, and O. B. Paulson. Generalizable patterns in neuroimaging: How many principal components? *Neuroimage*, 9(5):534–544, May 1999.
- [164] P. C. Hansen and D. P. O’Leary. The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing*, 14(6):1487–1503, Nov 1993.

- [165] D. L. Harrington, S. M. Rao, K. Y. Haaland, J. A. Bobholz, A. R. Mayer, J. R. Binderx, and R. W. Cox. Specialized neural systems underlying representations of sequential movements. *Journal of Cognitive Neuroscience*, 12(1):56–77, Jan 2000.
- [166] T. Hastie, R. Tibshirani, and J. H. Friedman. *The elements of statistical learning: Data mining, inference, and prediction*. Springer New York, 2001.
- [167] R. N. A. Henson, C. J. Price, M. D. Rugg, R. Turner, and K. J. Friston. Detecting latency differences in event-related BOLD responses: Application to words versus nonwords and initial versus repeated face presentations. *Neuroimage*, 15(1):83–97, Jan 2002.
- [168] J. Himberg, A. Hyvärinen, and F. Esposito. Validating the independent components of neuroimaging time series via clustering and visualization. *Neuroimage*, 22(3):1214–1222, Jul 2004.
- [169] W. S. Hinshaw. Image formation by nuclear magnetic resonance: The sensitive-point method. *Journal of Applied Physics*, 47(8):3709–3721, 1976.
- [170] A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, Feb 1970.
- [171] R. V. Hogg and A. Craig. *Introduction to Mathematical Statistics*. Prentice Hall, 5th edition, 1995.
- [172] A. M. Howseman and R. W. Bowtell. Functional magnetic resonance imaging: Imaging techniques and contrast mechanisms. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 354(1387):1179–1194, Jul 1999.
- [173] A.M. Howseman, M.K. Stehling, B. Chapman, R. Coxon, R. Turner, R.J. Ordridge, M.G. Cawley, P. Glover, P. Mansfield, and R.E. Coupland. Improvements in snap-shot nuclear magnetic resonance imaging. *British Journal of Radiology*, 61(729):822–828, Sep 1988.
- [174] X. Hu, T. H. Le, and K. Uğurbil. Evaluation of the early response in fMRI in individual subjects using short stimulus duration. *Magnetic Resonance in Medicine*, 37(6):877–884, Jun 1997.
- [175] C. M. Hurvich and C. L. Tsai. Regression and time series model selection in small samples. *Biometrika*, 76(2):297–307, Jun 1989.
- [176] R. A. Hutchinson, R. S. Niculescu, T. A. Keller, I. Rustandi, and T. M. Mitchell. Modeling fMRI data generated by overlapping cognitive processes with unknown onsets using Hidden Process Models. *Neuroimage*, 46(1):87–104, May 2009.
- [177] A. Hyvarinen. Fast and robust fixed-point algorithms for Independent Component Analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, May 1999.
- [178] W. Hänicke, K. D. Merboldt, D. Chien, M. L. Gyngell, H. Bruhn, and J. Frahm. Signal strength in subsecond FLASH magnetic resonance imaging: The dynamic approach to steady state. *Medical Physics*, 17(6):1004–1010, Nov-Dec 1990.

- [179] G. D. Jackson, A. Connelly, J. H. Cross, I. Gordon, and D. G. Gadian. Functional magnetic resonance imaging of focal seizures. *Neurology*, 44(5):850–856, May 1994.
- [180] G. M. James, P. Radchenko, and J. Lv. DASSO: Connections between the Dantzig Selector and LASSO. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 71(1):127–142, 2009.
- [181] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- [182] P. Jezzard and R. S. Balaban. Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine*, 34(1):65–73, Jul 1995.
- [183] P. Jezzard and S. Clare. Sources of distortion in functional MRI data. *Human Brain Mapping*, 8(2-3):80–85, 1999.
- [184] P. Jezzard, P. M. Matthews, and S. M. Smith. *Functional Magnetic Resonance Imaging: An Introduction to Methods*. Oxford University Press, 2001.
- [185] H. J. Jo, Z. S. Saad, W. K. Simmons, L. A. Milbury, and R. W. Cox. Mapping sources of correlation in resting state fMRI, with artifact detection and removal. *Neuroimage*, 52(2):571–582, Aug 2010.
- [186] O. Josephs, R. Turner, and K. J. Friston. Event-related fMRI. *Human Brain Mapping*, 5:243–248, May 1997.
- [187] R. Kawashima, P. E. Roland, and B. T. O’Sullivan. Activity in the human primary motor cortex related to ipsilateral hand movements. *Brain Research*, 663(2):251–256, Nov 1994.
- [188] K. N. Kay, S. V. David, R. J. Prenger, K. A. Hansen, and J. L. Gallant. Modeling low-frequency fluctuation and hemodynamic response timecourse in event-related fMRI. *Human Brain Mapping*, 29:142–156, Mar 2008.
- [189] S. M. Kay. *Fundamentals of Statistical Signal Processing: Detection Theory*. Prentice Hall, 1998.
- [190] P. Kellman and E. R. McVeigh. Image reconstruction in SNR units: A general method for SNR measurement. *Magnetic Resonance in Medicine*, 54(6):1439–1447, Dec 2005.
- [191] I. Khalidov, D. Van De Ville, M. J. Fadili, and M. Unser. Activelets and sparsity: A new way to detect brain activation from fMRI data. In *SPIE Wavelets XII*, 2007.
- [192] K. A. Kiehl, A. M. Smith, R. D. Hare, and P.F. Liddle. An event-related potential investigation of response inhibition in schizophrenia and psychopathy. *Biological Psychiatry*, 48(3):210–221, Aug 2000.
- [193] V. Kiviniemi, T. Starck, J. Remes, X. Long and J. Nikkinen, M. Haapea, J. Veijola, I. Moilanen, M. Isohanni, Y. F. Zang, and O. Tervonen. Functional segmentation of the brain cortex using high model order group PICA. *Human Brain Mapping*, 30(12):3865–3886, Dec 2009.

- [194] A. Kleinschmidt, C. Büchel, S. Zeki, and R. S. Frackowiak. Human brain activity during spontaneously reversing perception of ambiguous figures. *Proceedings of the Royal Society of London. Series B, Biological sciences*, 265(1413):2427–2433, Dec 1998.
- [195] A. Kleinschmidt, H. Obrig, M. Requardt, K. D. Merboldt, U. Dirnagl, A. Villringer, and J. Frahm. Simultaneous recording of cerebral blood oxygenation changes during human brain activation by magnetic resonance imaging and near-infrared spectroscopy. *Journal of Cerebral Blood Flow & Metabolism*, 16(5):817–826, Sep 1996.
- [196] E. Kobayashi, A. P. Bagshaw, C. Grova, F. Dubeau, and J. Gotman. Negative BOLD responses to epileptic spikes. *Human Brain Mapping*, 27(6):488–497, Jun 2006.
- [197] E. Kobayashi, C. S. Hawco, C. Grova, F. Dubeau, and J. Gotman. Widespread and intense BOLD changes during brief focal electrographic seizures. *Neurology*, 66(7):1049–1055, Apr 2006.
- [198] M. Kowalski. Sparse regression using mixed norms. *Applied and Computational Harmonic Analysis*, 27(3):303–324, Nov 2009.
- [199] F. Kruggel and D. Y. von Cramon. Modeling the hemodynamic response in single-trial functional MRI experiments. *Magnetic Resonance in Medicine*, 42(4):787–797, Oct 1999.
- [200] G. Krüger and G. H. Glover. Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magnetic Resonance in Medicine*, 46(4):631–637, Oct 2001.
- [201] G. Krüger, A. Kastrup, and G. H. Glover. Neuroimaging at 1.5 T and 3.0 T: Comparison of oxygenation-sensitive magnetic resonance imaging. *Magnetic Resonance in Medicine*, 45(4):595–604, Apr 2001.
- [202] A. Kumar, D. Welti, and R.R. Ernst. NMR Fourier zeugmatography. *Journal of Magnetic Resonance*, 18(1):69–83, Apr 1975.
- [203] K. K. Kwong, J. W. Belliveau, D. A. Chesler, I. E. Goldberg, R. M. Weisskoff, B. P. Poncelet, D. N. Kennedy, B. E. Hoppel, M. S. Cohen, and R. Turner. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences USA*, 89(12):5675–5679, Jun 1992.
- [204] S. M. LaConte, S. C. Ngan, and X. Hu. Wavelet transform-based Wiener filtering of event-related fMRI data. *Magnetic Resonance in Medicine*, 44(5):746–757, Nov 2000.
- [205] P.C. Lauterbur. Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature*, 242(5394):190–191, Mar 1973.
- [206] J. F. Lawless and P. Wang. A simulation study of ridge and other regression estimators. *Communications in Statistics-Theory and Methods*, 5(4):307–323, 1976.

- [207] N. Lazar. *The statistical analysis of functional MRI*. Springer, 2008.
- [208] K. Lee and J. C. Ye. Statistical parametric mapping of fMRI data using sparse dictionary learning. In *IEEE International Symposium on Biomedical Imaging, ISBI*, 2010.
- [209] L. Lemieux, A. Salek-Haddadi, T. E. Lund, H. Laufs, and D. Carmichael. Modelling large motion events in fMRI studies of patients with epilepsy. *Magnetic Resonance Imaging*, 25(6):894–901, Jul 2007.
- [210] P. LeVan, L. Tyvaert, F. Moeller, and J. Gotman. Independent component analysis reveals dynamic ictal BOLD responses in EEG-fMRI data from focal epilepsy patients. *Neuroimage*, 49(1):366–378, Jan 2010.
- [211] J. Li and P. Stoica. *Robust adaptive beamforming*. Wiley-Interscience, 2005.
- [212] Y. O. Li, T. Adali, and V. D. Calhoun. Estimating the number of independent components for functional magnetic resonance imaging data. *Human Brain Mapping*, 28(11):1251–1266, Nov 2007.
- [213] C. H. Liao, K. J. Worsley, J. B. Poline, J. A. D. Aston, G. H. Duncan, and A. C. Evans. Estimating the delay of the fMRI response. *Neuroimage*, 16(3 Pt 1):593–606, Jul 2002.
- [214] M. A. Lindquist, J. Meng-Loh, L. Y. Atlas, and T. D. Wager. Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling. *Neuroimage*, 45(1 Suppl):S187–S198, Mar 2009.
- [215] M. A. Lindquist, C. Waugh, and T. D. Wager. Modeling state-related fMRI activity using change-point theory. *Neuroimage*, 35(3):1125–1141, Apr 2007.
- [216] A. D. Liston, T. E. Lund, A. Salek-Haddadi, K. Hamandi, K. J. Friston, and L. Lemieux. Modelling cardiac signal as a confound in EEG-fMRI and its application in focal epilepsy studies. *Neuroimage*, 30(3):827–834, Apr 2006.
- [217] T. T. Liu. Efficiency, power, and entropy in event-related fMRI with multiple trial types. Part II: Design of experiments. *Neuroimage*, 21(1):401–413, Jan 2004.
- [218] T. T. Liu and L. R. Frank. Efficiency, power, and entropy in event-related fMRI with multiple trial types. Part I: Theory. *Neuroimage*, 21(1):387–400, Jan 2004.
- [219] T. T. Liu, L. R. Frank, E. C. Wong, and R. B. Buxton. Detection power, estimation efficiency, and predictability in event-related fMRI. *Neuroimage*, 13(4):759–773, Apr 2001.
- [220] Y. Liu, J. H. Gao, H. L. Liu, and P. T. Fox. The temporal response of the brain after eating revealed by functional MRI. *Nature*, 405(6790):1058–1062, Jun 2000.
- [221] S. Ljunggren. A simple graphical representation of Fourier-based imaging methods. *Journal of Magnetic Resonance*, 54(2):338–343, 1983.

- [222] N. K. Logothetis. What we can do and what we cannot do with fMRI. *Nature*, 453(7197):869–878, Jun 2008.
- [223] N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann. Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843):150–157, Jul 2001.
- [224] C. Long, E. N. Brown, D. Manoach, and V. Solo. Spatiotemporal wavelet analysis for functional MRI. *Neuroimage*, 23(2):500–516, Oct 2004.
- [225] C. J. Long, E. N. Brown, C. Triantafyllou, I. Aharon, L. L. Wald, and V. Solo. Nonstationary noise estimation in functional MRI. *Neuroimage*, 28(4):890–903, Dec 2005.
- [226] R. G. Lorenz and S. P. Boyd. Robust minimum variance beamforming. *IEEE Transactions Signal Processing*, 53(5):1684–1696, May 2005.
- [227] Y. Lu, T. Jiang, and Y. Zang. Single-trial variable model for event-related fMRI data analysis. *IEEE Transactions Medical Imaging*, 24(2):236–245, Feb 2005.
- [228] T. E. Lund and H. B. W. Larsson. Spatial distribution of low-frequency noise in fMRI. In *Proceedings of the 7th Annual Meeting of ISMRM*, page 1705, 1999.
- [229] T. E. Lund, K. H. Madsen, K. Sidaros, W. L. Luo, and T. E. Nichols. Non-white noise in fMRI: Does modelling have an impact? *Neuroimage*, 29(1):54–66, Jan 2006.
- [230] H. Luo and S. Puthusserypady. A sparse Bayesian method for determination of flexible design matrix for fMRI data analyses. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 12(12):2699–2706, Dec 2005.
- [231] H. Luo and S. Puthusserypady. Analysis of fMRI data with drift: Modified general linear model and Bayesian estimator. *IEEE Transactions on Biomedical Engineering*, 55(5):1504–1511, May 2008.
- [232] W. L. Luo and T. E. Nichols. Diagnosis and exploration of massively univariate neuroimaging models. *Neuroimage*, 19(3):1014–1032, Jul 2003.
- [233] S. Mallat. *A Wavelet Tour of Signal Processing: The Sparse Way, Third Edition*. Academic Press, 2008.
- [234] D. Malonek and A. Grinvald. Interactions between electrical activity and cortical microcirculation revealed by imaging spectroscopy: Implications for functional brain mapping. *Science*, 272(5261):551–554, Apr 1996.
- [235] P. Mansfield. Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 10(3):L55–L58, 1977.
- [236] P. Mansfield, R. Coxon, and J. Hykin. Echo-volumar imaging (EVI) of the brain at 3.0 T: First normal volunteer and functional imaging results. *Journal of Computer Assisted Tomography*, 19(6):847–852, Nov-Dec 1995.

- [237] P. Mansfield and P. K. Grannell. NMR 'diffraction' in solids? *Journal of Physics C: Solid State Physics*, 6(22):L422–L426, Nov 1973.
- [238] P. Mansfield and A.A. Maudsley. Medical imaging by NMR. *British Journal of Radiology*, 50:188–194, 1977.
- [239] J. L. Marini and S. M. Smith. On bias in the estimation of autocorrelations for fMRI voxel time-series analysis. *Neuroimage*, 18(1):83–90, Jan 2003.
- [240] J. P. Marques, J. Rebola, P. Figueiredo, A. Pinto, F. Sales, and M. Castelo-Branco. ICA decomposition of EEG signal for fMRI processing in epilepsy. *Human Brain Mapping*, 30(9):2986–2996, Sep 2009.
- [241] G. Marrelec, H. Benali, P. Ciuciu, M. Pelegrini-Issac, and J. B. Poline. Robust Bayesian estimation of the hemodynamic response function in event-related BOLD fMRI using basic physiological information. *Human Brain Mapping*, 19(1):1–17, May 2003.
- [242] D. J. McGonigle, A. M. Howseman, B. S. Athwal, K. J. Friston, R. S. Frackowiak, and A. P. Holmes. Variability in fMRI: An examination of intersession differences. *Neuroimage*, 11(6 Pt 1):708–734, Jun 2000.
- [243] M. McKeown, Y. J. Hu, and Z. J. Wang. ICA denoising for event-related fMRI studies. *Conf Proc IEEE Eng Med Biol Soc*, 1:157–161, 2005.
- [244] M. J. McKeown, T. P. Jung, S. Makeig, G. Brown, S. S. Kindermann, T. W. Lee, and T. J. Sejnowski. Spatially independent activity patterns in functional MRI data during the stroop color-naming task. *Proceedings of the National Academy of Sciences USA*, 95(3):803–810, Feb 1998.
- [245] M. J. McKeown, S. Makeig, G. G. Brown, T. P. Jung, S. S. Kindermann, A. J. Bell, and T. J. Sejnowski. Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping*, 6(3):160–188, 1998.
- [246] M. J. McKeown and T. J. Sejnowski. Independent component analysis of fMRI data: Examining the assumptions. *Human Brain Mapping*, 6(5-6):368–372, 1998.
- [247] G.C. McKinnon. Ultrafast interleaved gradient-echo-planar imaging on a standard scanner. *Magnetic Resonance in Medicine*, 30(5):609–616, Nov 1993.
- [248] Menon and Kim. Spatial and temporal limits in cognitive neuroimaging with fMRI. *Trends in Cognitive Sciences*, 3(6):207–216, Jun 1999.
- [249] R. S. Menon, D. C. Luknowsky, and J. S. Gati. Mental chronometry using latency-resolved functional MRI. *Proceedings of the National Academy of Sciences USA*, 95(18):10902–10907, Sep 1998.
- [250] R. S. Menon, S. Ogawa, X. Hu, J. P. Strupp, P. Anderson, and K. Uğurbil. BOLD based functional MRI at 4 Tesla includes a capillary bed contribution: echo-planar imaging correlates with previous optical imaging using intrinsic signals. *Magnetic Resonance in Medicine*, 33(3):453–459, Mar 1995.

- [251] F. G. Meyer. Wavelet-based estimation of a semiparametric generalized linear model of fMRI time-series. *IEEE Transactions on Medical Imaging*, 22(3):315–322, Mar 2003.
- [252] C. M. Michel, M. M. Murray, G. Lantz, S. Gonzalez, L. Spinelli, and R. Grave de Peralta. EEG source imaging. *Clinical Neurophysiology*, 115(10):2195–2222, Oct 2004.
- [253] K. J. Miller, M. denNijs, P. Shenoy, J. W. Miller, R. P. N. Rao, and J. G. Ojemann. Real-time functional brain mapping using electrocorticography. *Neuroimage*, 37(2):504–507, Aug 2007.
- [254] F. Moeller, P. Levan, and J. Gotman. Independent component analysis (ICA) of generalized spike wave discharges in fMRI: Comparison with general linear model-based EEG-fMRI. *Human Brain Mapping*, Mar 2010.
- [255] V. L. Morgan and J. C. Gore. Detection of irregular, transient fMRI activity in normal controls using 2dTCA: Comparison to event-related analysis using known timing. *Human Brain Mapping*, 30(10):3393–3405, Oct 2009.
- [256] V. L. Morgan, J. C. Gore, and B. Abou-Khalil. Cluster analysis detection of functional MRI activity in temporal lobe epilepsy. *Epilepsy Research*, 76(1):22–33, Aug 2007.
- [257] V. L. Morgan, Y. Li, B. Abou-Khalil, and J. C. Gore. Development of 2dTCA for the detection of irregular, transient BOLD activity. *Human Brain Mapping*, 29(1):57–69, Jan 2008.
- [258] V. L. Morgan, R. R. Price, A. Arain, P. Modur, and B. Abou-Khalil. Resting functional MRI with temporal clustering analysis for localization of epileptic activity without EEG. *Neuroimage*, 21(1):473–481, Jan 2004.
- [259] P. G. Morris. Magnetic resonance imaging and magnetic resonance spectroscopy assessment of brain function in experimental animals and man. *Journal of Psychopharmacology*, 13(4):330–336, Dec 1999.
- [260] J.P. Mugler and J.R. Brookeman. Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MPRAGE). *Magnetic Resonance in Medicine*, 15(1):152–157, Jul 1990.
- [261] K. Murphy, J. Bodurka, and P. A. Bandettini. How long to scan? The relationship between fMRI temporal signal to noise ratio and necessary scan duration. *Neuroimage*, 34(2):565–574, Jan 2007.
- [262] E. Neher, B. Sakmann, and J. H. Steinbach. The extracellular patch clamp: A method for resolving currents through individual open channels in biological membranes. *Pflügers Archiv - European Journal of Physiology*, 375(2):219–228, Jul 1978.
- [263] S. C. Ngan and X. Hu. Analysis of functional magnetic resonance imaging data using self-organizing mapping with spatial connectivity. *Magnetic Resonance in Medicine*, 41(5):939–946, May 1999.

- [264] R. K. Niazy, C. F. Beckmann, G. D. Iannetti, J. M. Brady, and S. M. Smith. Removal of fMRI environment artifacts from EEG data using optimal basis sets. *Neuroimage*, 28(3):720–737, Nov 2005.
- [265] T. Nichols and S. Hayasaka. Controlling the familywise error rate in functional neuroimaging: A comparative review. *Statistical Methods in Medical Research*, 12(5):419–446, Oct 2003.
- [266] Y. Nir, R. Mukamel, I. Dinstein, E. Privman, M. Harel, L. Fisch, H. Gelbard-Sagiv, S. Kipervasser, F. Andelman, M. Y. Neufeld, U. Kramer, A. Arieli, I. Fried, and R. Malach. Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. *Nature Neuroscience*, 11(9):1100–1108, Sep 2008.
- [267] P. L. Nunez and R. Srinivasan. *Electric fields of the brain: The neurophysics of EEG, 2nd Edition*. Oxford University Press, 2005.
- [268] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences USA*, 87(24):9868–9872, Dec 1990.
- [269] S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences USA*, 89(13):5951–5955, Jul 1992.
- [270] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab. *Signals & systems, 2nd Edition*. Prentice-Hall, Upper Saddle River, NJ, USA, 1996.
- [271] R. J. Ordidge, A. Howseman, R. Coxon, R. Turner, B. Chapman, P. Glover, M. Stehling, and P. Mansfield. Snapshot imaging at 0.5 T using echo-planar techniques. *Magnetic Resonance in Medicine*, 10(2):227–240, May 1989.
- [272] M. Osborne, B. Presnell, and B. A. Turlach. On the LASSO and its dual. *Journal of Computational and Graphical Statistics*, 9(2):319–337, Jun 2000.
- [273] D. Ostwald, C. Porcaro, and A. P. Bagshaw. An information theoretic approach to EEG-fMRI integration of visually evoked responses. *Neuroimage*, 49(1):498–516, Jan 2010.
- [274] R. M. Sanchez Panchuelo. *High resolution anatomical and functional imaging*. PhD thesis, University of Nottingham, 2009.
- [275] T. B. Parrish, D. R. Gitelman, K. S. LaBar, and M. M. Mesulam. Impact of signal-to-noise on functional MRI. *Magnetic Resonance in Medicine*, 44(6):925–932, Dec 2000.
- [276] A. Pascual-Leone, D. Bartres-Faz, and J. P. Keenan. Transcranial magnetic stimulation: Studying the brain-behaviour relationship by induction of 'virtual lesions'. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 354(1387):1229–1238, Jul 1999.
- [277] A. Pascual-Leone, N. J. Davey, J. Rothwell, E. M. Wassermann, and B. K. Puri. *Handbook of transcranial magnetic stimulation*. Oxford, 2002.

- [278] R. E. Passingham. Premotor cortex: Sensory cues and movement. *Behavioural Brain Research*, 18(2):175–185, Nov-Dec 1985.
- [279] L. Pauling and C. D. Coryell. The magnetic properties and structure of hemoglobin, oxyhemoglobin and carbonmonoxyhemoglobin. *Proceedings of the National Academy of Sciences USA*, 22(4):210–216, Apr 1936.
- [280] W. Penny, S. Kiebel, and K. Friston. Variational Bayesian inference for fMRI time series. *Neuroimage*, 19(3):727–741, Jul 2003.
- [281] W. D. Penny, N. J. Trujillo-Barreto, and K. J. Friston. Bayesian fMRI time series analysis with spatial priors. *Neuroimage*, 24(2):350–362, Jan 2005.
- [282] V. Perlberg, P. Bellec, J. L. Anton, M. Pélégriani-Issac, J. Doyon, and H. Benali. CORSICA: Correction of structured noise in fMRI by automatic identification of ICA components. *Magnetic Resonance Imaging*, 25(1):35–46, Jan 2007.
- [283] A. M. Peters, M. J. Brookes, F. G. Hoogenraad, P. A. Gowland, S. T. Francis, P. G. Morris, and R. Bowtell. T2* measurements in human brain at 1.5, 3 and 7 T. *Magnetic Resonance Imaging*, 25(6):748–753, Jul 2007.
- [284] N. Petridou, C. Caballero, I. Dryden, S. Francis, and P. Gowland. Spatiotemporally coordinated activation detected during apparent rest in fMRI. *NeuroImage*, 47(S1):S82, 2009.
- [285] J. Pfeuffer, J. C. McCullough, P. F. Van de Moortele, K. Uğurbil, and X. Hu. Spatial dependence of the nonlinear BOLD response at short stimulus duration. *Neuroimage*, 18(4):990–1000, Apr 2003.
- [286] J. Pfeuffer, P. F. van de Moortele, E. Yacoub, A. Shmuel, G. Adriany, P. Andersen, H. Merkle, M. Garwood, K. Uğurbil, and X. Hu. Zoomed functional imaging in the human brain at 7 Tesla with simultaneous high spatial and high temporal resolution. *Neuroimage*, 17(1):272–286, Sep 2002.
- [287] A. Pitiot, J. Totman, and P. Gowland. Null point imaging: A joint acquisition analysis paradigm for MR classification. *Medical Image Computing and Computer Assisted Intervention, MICCAI*, 10:759–766, 2007.
- [288] M. Poole and R. Bowtell. Volume parcellation for improved dynamic shimming. *MAGMA*, 21(1-2):31–40, Mar 2008.
- [289] C. A. Porro, V. Cettolo, M. P. Francescato, and P. Baraldi. Ipsilateral involvement of primary motor cortex during motor imagery. *European Journal of Neuroscience*, 12(8):3059–3063, Aug 2000.
- [290] C. A. Porro, M. P. Francescato, V. Cettolo, M. E. Diamond, P. Baraldi, C. Zuiani, M. Bazzocchi, and P. E. di Prampero. Primary motor and sensory cortex activation during motor performance and motor imagery: A functional magnetic resonance imaging study. *Journal of Neuroscience*, 16(23):7688–7698, Dec 1996.
- [291] B. A. Poser and D. G. Norris. Investigating the benefits of multi-echo EPI for fMRI at 7 T. *Neuroimage*, 45(4):1162–1172, May 2009.

- [292] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962, Nov 1999.
- [293] E. M. Purcell, H. C. Torrey, and R. V. Pound. Resonance absorption by nuclear magnetic moments in a solid. *Physical Review*, 69(1-2):37–38, Jan 1946.
- [294] I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch. A new method of measuring nuclear magnetic moment. *Physical Review*, 53(4):318, Feb 1938.
- [295] J. C. Rajapakse, F. Kruggel, J. M. Maisog, and D. Y. von Cramon. Modeling hemodynamic response for analysis of functional MRI time-series. *Human Brain Mapping*, 6(4):283–300, 1998.
- [296] S. M. Rao, J. R. Binder, P. A. Bandettini, T. A. Hammeke, F. Z. Yetkin, A. Jesmanowicz, L. M. Lisk, G. L. Morris, W. M. Mueller, and L. D. Estkowski. Functional magnetic resonance imaging of complex human movements. *Neurology*, 43(11):2311–2318, Nov 1993.
- [297] M. Razavi, T. J. Grabowski, W. P. Vispoel, Patrick Monahan, S. Mehta, B. Eaton, and L. Bolinger. Model assessment and model building in fMRI. *Human Brain Mapping*, 20(4):227–238, Dec 2003.
- [298] W. Richter, P. M. Andersen, A. P. Georgopoulos, and S. G. Kim. Sequential activity in human motor areas during a delayed cued finger movement task studied by time-resolved fMRI. *Neuroreport*, 8(5):1257–1261, Mar 1997.
- [299] W. Richter, R. Somorjai, R. Summers, M. Jarmasz, R. S. Menon, J. S. Gati, A. P. Georgopoulos, C. Tegeler, K. Uğurbil, and S. G. Kim. Motor area activity during mental rotation studied by time-resolved single-trial fMRI. *Journal of Cognitive Neuroscience*, 12(2):310–320, Mar 2000.
- [300] W. Richter, K. Uğurbil, A. Georgopoulos, and S. G. Kim. Time-resolved fMRI of mental rotation. *Neuroreport*, 8(17):3697–3702, Dec 1997.
- [301] F. De Ridder, R. Pintelon, J. Schoukens, , and D.P. Gillikin. Modified aic and mdl model selection criteria for short data records. *IEEE Transactions on Instrumentation and Measurement*, 54(1):144–150, Feb 2005.
- [302] J. Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, Sep 1978.
- [303] L. F. Robinson, T. D. Wager, and M. A. Lindquist. Change point estimation in multi-subject fMRI studies. *Neuroimage*, 49(2):1581–1592, Jan 2010.
- [304] R. Rodionov, F. De Martino, H. Laufs, D. W. Carmichael, E. Formisano, M. Walker, John S Duncan, and L. Lemieux. Independent component analysis of interictal fMRI in focal epilepsy: Comparison with general linear model-based EEG-correlated fMRI. *Neuroimage*, 38(3):488–500, Nov 2007.
- [305] C. S. Roy and C. S. Sherrington. On the regulation of the blood-supply of the brain. *Journal of Physiology*, 11(1-2):85–158.17, Jan 1890.
- [306] R. Rubinstein, A. M. Bruckstein, and M. Elad. Dictionaries for sparse representation modeling. *Proceedings of the IEEE*, 98(6):1045 –1057, Jun 2010.

- [307] S. Ryali, K. Supekar, D. A. Abrams, and V. Menon. Sparse logistic regression for whole-brain classification of fMRI data. *Neuroimage*, 51(2):752–764, Jun 2010.
- [308] A. Salek-Haddadi, K. J. Friston, L. Lemieux, and D. R. Fish. Studying spontaneous EEG activity with fMRI. *Brain Research Reviews*, 43(1):110–133, Sep 2003.
- [309] F. Di Salle, E. Formisano, E. Seifritz, D. E. Linden, K. Scheffler, C. Saulino, G. Tedeschi, F. E. Zanella, A. Pepino, R. Goebel, and E. Marciano. Functional fields in human auditory cortex revealed by time-resolved fMRI without interference of EPI noise. *Neuroimage*, 13(2):328–338, Feb 2001.
- [310] S. Sardy, A. Bruce, and P. Tseng. Block coordinate relaxation methods for nonparametric wavelet denoising. *Journal of Computational and Graphical Statistics*, 9(2):361–379, 2000.
- [311] G.E. Sarty. *Computing brain activity maps from fMRI time-series images*. Cambridge University Press, 2007.
- [312] F. E. Satterthwaite. An approximate distribution of estimates of variance components. *Biometrics Bulletin*, 2(6):110–114, 1946.
- [313] T. Schubert, D. Y. von Cramon, T. Niendorf, S. Pollmann, and P. Bublak. Cortical areas and the control of self-determined finger movements: An fMRI study. *Neuroreport*, 9(14):3171–3176, Oct 1998.
- [314] A. Schwartzman, R. F. Dougherty, J. Lee, D. Ghahremani, and J. E. Taylor. Empirical null and false discovery rate analysis in neuroimaging. *Neuroimage*, 44(1):71–82, Jan 2009.
- [315] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, Mar 1978.
- [316] C. Schwarzbauer, T. Mildner, W. Heinke, M. Brett, and R. Deichmann. Dual echo EPI—the method of choice for fMRI in the presence of magnetic field inhomogeneities? *Neuroimage*, 49(1):316–326, Jan 2010.
- [317] Y. Selen, R. Abrahamsson, and P. Stoica. Automatic robust adaptive beamforming via ridge regression. *Signal Processing*, 88(1):33 – 49, Jan 2008.
- [318] I. W. Selesnick and I. Bayram. Oscillatory + transient signal decomposition using overcomplete rational-dilation wavelet transforms. In *Proceedings of SPIE*, volume 7446, 2009.
- [319] A. Shmuel, M. Augath, A. Oeltermann, and N. K Logothetis. Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nature Neuroscience*, 9(4):569–577, Apr 2006.
- [320] K. Shmueli, P. van Gelderen, J. A. de Zwart, S. G. Horovitz, M. Fukunaga, J. M. Jansma, and J. H. Duyn. Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal. *Neuroimage*, 38(2):306–320, Nov 2007.

- [321] A. M. Smith, B. K. Lewis, U. E. Ruttimann, F. Q. Ye, T. M. Sinnwell, Y. Yang, J. H. Duyn, and J. A. Frank. Investigation of low frequency drift in fMRI signal. *Neuroimage*, 9(5):526–533, May 1999.
- [322] S. Smith, M. Jenkinson, C. Beckmann, K. Miller, and M. Woolrich. Meaningful design and contrast estimability in fMRI. *Neuroimage*, 34(1):127–136, Jan 2007.
- [323] S. M. Smith, C. F. Beckmann, N. Ramnani, M. W. Woolrich, P. R. Bannister, M. Jenkinson, P. M. Matthews, and D. J. McGonigle. Variability in fMRI: A re-examination of inter-session differences. *Human Brain Mapping*, 24(3):248–257, Mar 2005.
- [324] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. J. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, and P. M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, 23(S1):S208–S219, 2004.
- [325] S. M. Smith and T. E. Nichols. Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage*, 44(1):83–98, Jan 2009.
- [326] A. Smolders, F. De Martino, N. Staeren, P. Scheunders, J. Sijbers, R. Goebel, and E. Formisano. Dissecting cognitive stages with time-resolved fMRI data: A comparison of Fuzzy Clustering and Independent Component Analysis. *Magnetic Resonance Imaging*, 25(6):860–868, Jul 2007.
- [327] D. A. Soltysik, K. K. Peck, K. D. White, B. Crosson, and R. W. Briggs. Comparison of hemodynamic response nonlinearity across primary cortical areas. *Neuroimage*, 22(3):1117–1127, Jul 2004.
- [328] A.W. Song, E.C. Wong, and J.S. Hyde. Echo-volume imaging. *Magnetic Resonance in Medicine*, 32(5):668–671, Nov 1994.
- [329] O. Speck, J. Stadler, and M. Zaitsev. High resolution single-shot EPI at 7t. *MAGMA*, 21(1-2):73–86, Mar 2008.
- [330] J. L. Starck, M. Elad, and D. L. Donoho. Redundant multiscale transforms and their applications for morphological component analysis. *Advances in Imaging and Electron Physics*, 132:287–348, 2004.
- [331] J. L. Starck, M. Elad, and D. L. Donoho. Image decomposition via the combination of sparse representations and a variational approach. *IEEE Transactions on Image Processing*, 14(10):1570–1582, Oct 2005.
- [332] C. M. Stein. Estimation of the mean of a multivariate Normal distribution. *The Annals of Statistics*, 9(6):1135–1151, Nov 1981.
- [333] P. Stoica and R. L. Moses. *Spectral Analysis of Signals*. Prentice-Hall, Upper Saddle River, NJ, 2005.
- [334] S. C. Strother. Evaluating fMRI preprocessing pipelines. *IEEE Engineering in Medicine and Biology Magazine*, 25(2):27–41, Mar-Apr 2006.

- [335] M. Svensén, F. Kruggel, and D. Y. von Cramon. Probabilistic modeling of single-trial fMRI data. *IEEE Transactions Medical Imaging*, 19(1):25–35, Jan 2000.
- [336] J. Tanabe, D. Miller, J. Tregellas, R. Freedman, and F. G. Meyer. Comparison of detrending methods for optimal fMRI preprocessing. *Neuroimage*, 15(4):902–907, Apr 2002.
- [337] B. Thirion and O. Faugeras. Dynamical components analysis of fMRI data through kernel PCA. *Neuroimage*, 20(1):34–49, Sep 2003.
- [338] K. R. Thulborn, J. C. Waterton, P. M. Matthews, and G. K. Radda. Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochimica et Biophysica Acta*, 714(2):265–270, Feb 1982.
- [339] R. Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 58(1):267–288, 1996.
- [340] M. E. Tipping. Sparse bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244, Sep 2001.
- [341] F. Tong, M. Meng, and R. Blake. Neural bases of binocular rivalry. *Trends in Cognitive Sciences*, 10(11):502–511, Nov 2006.
- [342] C. Triantafyllou, M. Elschot, J. Polimeni, and L. Wald. Physiological noise in GRAPPA fMRI time-series. In *17th Scientific Meeting International Society for Magnetic Resonance in Medicine (ISMRM)*, 2009.
- [343] C. Triantafyllou, R. D. Hoge, G. Krueger, C. J. Wiggins, A. Potthast, G. C. Wiggins, and L. L. Wald. Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage*, 26(1):243–250, May 2005.
- [344] L. Trojano, D. Grossi, D. E. Linden, E. Formisano, H. Hacker, F. E. Zanella, R. Goebel, and F. Di Salle. Matching two imagined clocks: The functional anatomy of spatial analysis in the absence of visual stimulation. *Cerebral Cortex*, 10(5):473–481, May 2000.
- [345] J. A. Tropp and S. J. Wright. Computational methods for sparse solution of linear inverse problems. *Proceedings of the IEEE*, 98(6):948–958, Jun 2010.
- [346] R. Turner, D. Le Bihan, C. T. Moonen, D. Despres, and J. Frank. Echo-planar time course MRI of cat brain oxygenation changes. *Magnetic Resonance in Medicine*, 22(1):159–166, Nov 1991.
- [347] D. B. Twieg. The k-trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods. *Medical Physics*, 10(5):610–621, Sep 1983.
- [348] L. Tyvaert, C. Hawco, E. Kobayashi, P. LeVan, F. Dubeau, and J. Gotman. Different structures involved during ictal and interictal epileptic activity in malformations of cortical development: An EEG-fMRI study. *Brain*, 131(8):2042–2060, Aug 2008.

- [349] K. Uğurbil, X. Hu, W. Chen, X. H. Zhu, S. G. Kim, and A. Georgopoulos. Functional mapping in the human brain using high magnetic fields. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 354(1387):1195–1213, Jul 1999.
- [350] W. van der Zwaag, S. Francis, K. Head, A. Peters, P. Gowland, P. Morris, and R. Bowtell. fMRI at 1.5, 3 and 7 T: Characterising BOLD signal changes. *Neuroimage*, 47(4):1425–1434, Oct 2009.
- [351] P. J. van Houdt, P. P. W. Ossenblok, P. A. J. M. Boon, F. S. S. Leijten, D. N. Velis, C. J. Stam, and J. C. de Munck. Correction for pulse height variability reduces physiological noise in functional MRI when studying spontaneous brain activity. *Human Brain Mapping*, 31(2):311–325, Feb 2010.
- [352] A. F. van Rootselaar, R. Renken, B. M. de Jong, J. M. Hoogduin, M. A. J. Tijssen, and N. M. Maurits. fMRI analysis for motor paradigms using EMG-based designs: A validation study. *Human Brain Mapping*, 28(11):1117–1127, Nov 2007.
- [353] A. L. Vazquez and D. C. Noll. Nonlinear aspects of the BOLD response in functional MRI. *Neuroimage*, 7(2):108–118, Feb 1998.
- [354] D. Van De Ville, M. L. Seghier, F. Lazeyras, T. Blu, and M. Unser. Wspm: Wavelet-based statistical parametric mapping. *Neuroimage*, 37(4):1205–1217, Oct 2007.
- [355] A. Villringer, J. Planck, C. Hock, L. Schleinkofer, and U. Dirnagl. Near infrared spectroscopy (NIRS): A new tool to study hemodynamic changes during activation of brain function in human adults. *Neuroscience Letters*, 154(1-2):101–104, May 1993.
- [356] P. Vuilleumier, R. N. Henson, J. Driver, and R. J. Dolan. Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience*, 5:491–499, 2002.
- [357] S. Vulliemoz, L. Lemieux, J. Daunizeau, C. M. Michel, and J. S. Duncan. The combination of EEG source imaging and EEG-correlated functional MRI to map epileptic networks. *Epilepsia*, 51(4):491–505, Apr 2010.
- [358] S. Vulliemoz, R. Rodionov, D. W. Carmichael, R. Thornton, M. Guye, S. D. Lhatoo, C. M. Michel, J. S. Duncan, and L. Lemieux. Continuous EEG source imaging enhances analysis of EEG-fMRI in focal epilepsy. *Neuroimage*, 49(4):3219–3229, Feb 2010.
- [359] S. Vulliemoz, R. Thornton, R. Rodionov, D. W. Carmichael, M. Guye, S. Lhatoo, A. W. McEvoy, L. Spinelli, C. M. Michel, J. S. Duncan, and L. Lemieux. The spatio-temporal mapping of epileptic networks: Combination of EEG-fMRI and EEG source imaging. *Neuroimage*, 46(3):834–843, Jul 2009.
- [360] T. D. Wager, A. Vazquez, L. Hernandez, and D. C. Noll. Accounting for nonlinear BOLD effects in fMRI: Parameter estimates and a model for prediction in rapid event-related studies. *Neuroimage*, 25(1):206–218, Mar 2005.

- [361] N. Weiskopf, C. Hutton, O. Josephs, and R. Deichmann. Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: A whole-brain analysis at 3 T and 1.5 T. *Neuroimage*, 33(2):493–504, Nov 2006.
- [362] R. M. Weisskoff, J. Baker, J. Belliveau, T.L. Davis, K.K. Kwong, M. Cohen, and B. R. Rosen. Power spectrum analysis of functionally weighted MR data: What’s in the noise? In *Proceedings SMRM*, volume 12, page 7, 1993.
- [363] C. Windischberger, R. Cunnington, C. Lamm, R. Lanzenberger, H. Langenberger, L. Deecke, H. Bauer, and E. Moser. Time-resolved analysis of fMRI signal changes using Brain Activation Movies. *Journal of Neuroscience Methods*, 169(1):222–230, Mar 2008.
- [364] C. Windischberger, C. Lamm, H. Bauer, and E. Moser. Consistency of inter-trial activation using single-trial fMRI: Assessment of regional differences. *Cognitive Brain Research*, 13(1):129–138, Feb 2002.
- [365] C. Windischberger, H. Langenberger, T. Sycha, E. M. Tschernko, G. Fuchsjaeger-Mayerl, L. Schmetterer, and E. Moser. On the origin of respiratory artifacts in BOLD-EPI of the human brain. *Magnetic Resonance Imaging*, 20(8):575–582, Oct 2002.
- [366] D. Wipf and S. Nagarajan. Iterative reweighted ell_1 and ell_2 methods for finding sparse solutions. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):317–329, Apr 2010.
- [367] D. P. Wipf and B. D. Rao. Sparse Bayesian learning for basis selection. *IEEE Transactions on Signal Processing*, 52(8):2153–2164, Aug. 2004.
- [368] S. T. Witt, A. R. Laird, and M. E. Meyerand. Functional neuroimaging correlates of finger-tapping task variations: An ALE meta-analysis. *NeuroImage*, 42(1):343–356, Aug 2008.
- [369] M. W. Woolrich, B. D. Ripley, M. Brady, and S. M. Smith. Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage*, 14(6):1370–1386, Dec 2001.
- [370] K. J. Worsley. Detecting activation in fMRI data. *Statistical Methods in Medical Research*, 12(5):401–418, Oct 2003.
- [371] K. J. Worsley. Spatial smoothing of autocorrelations to control the degrees of freedom in fMRI analysis. *Neuroimage*, 26(2):635–641, Jun 2005.
- [372] K. J. Worsley and K. J. Friston. Analysis of fMRI time-series revisited - again. *Neuroimage*, 2(3):173–181, Sep 1995.
- [373] K. J. Worsley, C. H. Liao, J. Aston, V. Petre, G. H. Duncan, F. Morales, and A. C. Evans. A general statistical analysis for fMRI data. *Neuroimage*, 15(1):1–15, Jan 2002.
- [374] K. J. Worsley and J. E. Taylor. Detecting fMRI activation allowing for unknown latency of the hemodynamic response. *Neuroimage*, 29(2):649–654, Jan 2006.

- [375] P. J. Wright, O. E. Mougin, J. J. Totman, A. M. Peters, M. J. Brookes, R. Coxon, P. E. Morris, M. Clemence, S. T. Francis, R. W. Bowtell, and P. A. Gowland. Water proton T1 measurements in brain tissue at 7, 3, and 1.5 T using IR-EPI, IR-TSE, and MPRAGE: Results and optimization. *MAGMA*, 21(1-2):121–130, Mar 2008.
- [376] E. Yacoub, T. Q. Duong, P. F. Van De Moortele, M. Lindquist, G. Adriany, S. G. Kim, K. Uğurbil, and X. Hu. Spin-echo fMRI in humans using high spatial resolutions and high magnetic fields. *Magnetic Resonance in Medicine*, 49(4):655–664, Apr 2003.
- [377] E. Yacoub, N. Harel, and K. Uğurbil. High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences USA*, 105(30):10607–10612, Jul 2008.
- [378] E. Yacoub, T. H. Le, K. Uğurbil, and X. Hu. Further evaluation of the initial negative response in functional magnetic resonance imaging. *Magnetic Resonance in Medicine*, 41(3):436–441, Mar 1999.
- [379] O. Yamashita, M. Sato, T. Yoshioka, F. Tong, and Y. Kamitani. Sparse estimation automatically selects voxels relevant for the decoding of fMRI activity patterns. *Neuroimage*, 42(4):1414–1429, Oct 2008.
- [380] D. Yekutieli and Y. Benjamini. Resampling-based false discovery rate controlling multiple test procedures for correlated test statistics. *Journal of Statistical Planning and Inference*, 82(1-2):171 – 196, Dec 1999.
- [381] E. Zarahn, G. Aguirre, and M. D’Esposito. A trial-based experimental design for fMRI. *Neuroimage*, 6(2):122–138, Aug 1997.
- [382] E. Zarahn, G. K. Aguirre, and M. D’Esposito. Empirical analyses of BOLD fMRI statistics. I. spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage*, 5(3):179–197, Apr 1997.
- [383] Y. Zhou and H. Liang. Statistical inference for semiparametric varying-coefficient partially linear models with error-prone linear covariates. *The Annals of Statistics*, 37(1):427–458, Feb 2009.
- [384] H. Zou, T. Hastie, and R. Tibshirani. On the degrees of freedom of the LASSO. *The Annals of Statistics*, 35(5):2173–2192, Oct 2007.