Applications of realtime-fMRI for non-invasive brain computer interfaces - decoding and neurofeedback

Jinendra Ekanayake
Supervisor: Professor Geraint Rees

Dissertation submitted
for the degree of Doctor of Philosophy,
Wellcome Trust Centre for Neuroimaging,
University College London
2015

‘ I, **Jinendra Ekanayake** confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.'

Acknowledgements

Undertaking this has been a life-changing experience on every possible level. As with all the best things in my life, it happened without much fanfare, a quiet celebration shared with a few truly remarkable individuals.

Thank you **Geraint Rees**. It is not possible to put into words how inspirational, and supportive you have been. I am where I am because of you. I will always be grateful to you. May my actions speak louder than these words.

Thank you **Eva Feredoes**. I love you. Not just for the support, and clear-sighted guidance through the last few years, but also for making this journey a happy and meaningful one, through all of the challenges. The best finding of my thesis without a shadow of doubt.

Thank you **Joel Winston**. Clever insights, invaluable scripting, and punishing wit. What more do you need!

The Wellcome Trust for Neuroimaging was an important place for me, introducing me to a community of like-minded individuals. I would especially like to thank **Eric Featherstone** who helped me to make the painful transition from clinical medicine to research programing. **Sheila Burns**, who helped me scan the Allegra out of the building (literally). **Peter Zeidman** for all kinds of help, and information (academic and otherwise). **Ged Ridgway** for making the realtime dream a reality, and putting up with my endless questions. **Nikolaus Weiskopf**, for making me ask the important questions, and the casual rescue of experiments, and failing scanners. **Chloe Hutton** for patiently putting up with an ignorant doctor. **Frank Scharnowski**, my research big brother, sharer of scripts, and all-round good guy.

I would like to thank my volunteers, who showed commitment and dedication to my studies (as evidenced by being scanned 6-7 days in a row, putting up with induced double vision, pieces of cardboard stuck to their foreheads and the temperamental but loveable Allegra!).

I would like to thank my clinical supervisors and mentors who allowed me to have two lives. **Neil Kitchen** for your support and guidance. **Laurence Watkins** for your tireless patience, and quiet confidence. **Ludvic Zrinzo** for giving me the time and raw materials to take the next step.

An honourable mention goes to **Haji Cham** ‘if a dog bites you, do you bite back ?’.

My deepest thank you and love to my parents and my grandmother. For the grit, laughter, and belief, but most of all for the unconditional love. When the pressure was on, we held. We are a small and indomitable team- watch out world, part 2 is about to begin!

Thesis Abstract

Non-invasive brain-computer interfaces (BCIs) seek to enable or restore brain function by using neuroimaging e.g. functional magnetic resonance imaging (fMRI), to engage brain activations without the need for explicit behavioural output or surgical implants. Brain activations are converted into output signals, for use in communication interfaces, motor prosthetics, or to directly shape brain function via a feedback loop.

The aim of this thesis was to develop cognitive BCIs using realtime fMRI (rt-fMRI), with the potential for use as a communication interface, or for initiating neural plasticity to facilitate neurorehabilitation. Rt-fMRI enables brain activation to be manipulated directly to produce changes in function, such as perception.

Univariate and multivariate classification approaches were used to decode brain activations produced by the deployment of covert spatial attention to simple visual stimuli. Primary and higher order visual areas were examined, as well as potential control regions. The classification platform was then developed to include the use of real-world visual stimuli, exploiting the use of category-specific visual areas, and demonstrating real-world applicability as a communications interface. Online univariate classification of spatial attention was successfully achieved, with individual classification accuracies for 4-quadrant spatial attention reaching 70%. Further, a novel implementation of m-sequences enabled the use of the timing of stimuli presentation to enhance signal characterisation.

An established rt-fMRI analysis loop was then used for neurofeedback-led manipulation of category-specific visual brain regions, modulating their functioning, and, as a result, biasing visual perception during binocular rivalry. These changes were linked with functional and effective connectivity changes in trained regions, as well as in a putative top-down control region.

The work presented provides proof-of-principle for non-invasive BCIs using rt-fMRI, with the potential for translation into the clinical environment. Decoding and

neurofeedback applied to non-invasive and implantable BCIs form an evolving continuum of options for enabling and restoring brain function.

Key words:

Realtime fMRI (rt-fMRI), Non-invasive brain computer interfaces (BCI), decoding, neurofeedback

Table of Contents

Table of Figures

Table of Tables

Table of Abbreviations

ACC	Anterior cingulate cortex
ANOVA	Analysis of Variance
BCI	Brain-computer interface
BOLD	Blood oxygen level dependent signal
CBF	Cerebral blood flow
CBV	Cerebral blood volume
$CMRO_2$	Cerebral metabolic rate for oxygen
CVA	Canonical Variate Analysis
DARTEL	Diffeomorphic Anatomical Registration Through Exponential Lie Algebra
DBS	Deep Brain Stimulation
DCM	Dynamic Causal Modelling
DLPFC	Dorsolateral prefrontal cortex
DTI	Diffusion tensor imaging
EBA	Extrastriate body area
ECoG	Electrocorticography
EEG	Electroencephalography
EPI	Echo Planar Imaging
FDR	False discovery rate
FDR	False Discovery Rate
FFA	Fusiform Face Area
FHWH	Full-width-at-half-maximum
FID	Free induction decay
fMRI	Functional magnetic resonance imaging
FWE	Familywise error
GABA	Gamma- aminobutyric acid
GE	Gradient echo
GLM	General linear model
HD	Huntington's disease
HRF	Haemodynamic response function
LDA	Linear Discriminant Analysis

LFP	Local field potentials
LOC	Lateral Occipital Cortex
MEG	Magnetoencephalography
MNI	Montreal Neurological Institute (convention for canonical brain space)
MRI	Magnetic Resonance Imaging
MUA	Multiple-unit spiking activity
MVPA	Multivariate pattern analysis
NFB	Neurofeedback
OFA	Occipital face area
PCA	Principal component analysis
PCC	Posterior cingulate cortex
PD	Parkinson’s disease
PPA	Parahippocampal Place Area
PSC	Percentage signal change
RF	Radiofrequency
RFT	Random Field theory
ROI	Region of interest
rt-fMRI	Realtime functional magnetic resonance imaging
SD	Standard deviation
SE	Spin echo
SEM	Standard error of the mean
SMA	Supplementary motor area
SPM	Statistical Parametric Mapping
SVM	Support vector machine
TBM	Tensor Based Morphometry
TBV	Turbo brain voyager analysis package
TE	Echo time
TR	Repetition time

PREFACE

Themes and Concepts

The aim of this thesis is to investigate the basis of a cognitive brain-computer-interface (BCI) using realtime functional magnetic resonance imaging (rt-fMRI), that could be developed for patient benefit either as a prosthesis for communication, or as a means of engaging neural plasticity more directly for the purposes of neurorehabilitation.

In order to achieve this goal, I attempt to answer two related but distinct questions. The first is technical - can a cutting edge technique, rt-fMRI, be implemented for the purposes of decoding and neurofeedback training, utilising higher order visual areas?

The second question is biological- can rt-fMRI be used to manipulate cognitive processes in real time, and thus gain mechanistic and causal insights into human brain function? Rt-fMRI enables brain activation to be manipulated directly; brain activation can be treated as an independent variable to produce changes in brain function, such as perception.

In the first three experimental chapters, I examine the decoding of the direction of spatial attention from higher order visual cortex, using both univariate and multivariate approaches, leading to the implementation of a rt-fMRI framework for the detection of covert spatial attention from brain activity. In the fourth chapter, I examine neurofeedback training of category-specific visual areas, and its effect on structural and functional plasticity. I set out to test whether high-level visual perception can be selectively biased by neurofeedback training. The studies in this thesis provide proof-of-principle of a non-invasive BCI using rt-fMRI, with the potential for translation into the clinical environment. In the final chapter I conclude with a discussion of developing the findings in this thesis for applications in the clinical environment.

1. GENERAL INTRODUCTION

1.1 Theme 1

Developing a cognitive brain-computer interface

1.1.1 Introduction

Cognitive BCIs aim to provide control of external devices or computers, using physiological signals related to brain activation only. The basis of the physiological signal can be acquired invasively, such as through implanted electrodes which record brain activity, using electrocorticography (ECoG), or noninvasively using imaging modalities such as functional magnetic resonance imaging (fMRI), electroencephalography (EEG) and magnetoencephalography (MEG) (1–9). The latter has the advantage of not requiring surgery, and its attendant risks, and may in fact be used to optimise surgical targets for neural interfaces.

BCI technology for restoration of function in patients uses two principal types of neural signal – 1) sensorimotor output for rehabilitation or control of a motor prosthesis, or 2) cognitive output from high order sensory cortex for communication and perceptual enhancement. This thesis is exclusively concerned with the latter approach, using signals associated with cognitive processes in higher order visual areas with the non-invasive modality of rt-fMRI.

1.1.2 BCIs- an overview

Rt-fMRI was selected as the basis for the development of a non-invasive BCIs due to a specific combination of attributes. FMRI provides superior spatial specificity and resolution, as compared to other comparable non-invasive imaging modalities e.g. MEG/EEG based devices (2). Functional brain regions can be identified, and

investigated in relation to task-based activations, and can be used to hone activation profiles, to increase signal strength, and improve decoding accuracy (10). The selection of an imaging technique such as fMRI for this thesis, which enables a fine-grained mapping of anatomical structure to function, was motivated by a long-term investment in BCI development. To begin, cortical structures at a high degree of spatial resolution, can be examined. Subsequently, the neurobiology of the region can be explored, enabling more complex assessments such as network connectivity between cortical and subcortical brain areas to be investigated (11,12). This is of particular importance in the eventual clinical translation of the work presented here, in diseases such as Parkinson's disease (PD) and Huntington's disease (HD). These disease processes affect basal ganglia structures, and may be addressed by targeting cortical structures such as the supplementary motor area (SMA) for neurofeedback based therapies, causing linked network changes in deep brain structures such as the subthalamic nucleus (STN) (13). The use of computer software which enables the rapid analysis of fMRI data in relation to task-related brain activation, such that it can be acted upon while the experiment is progress, allows an interactive facility for the patient or participant being scanned. This in particular allows brain activity to be used in relation to a set goal to produce a surrogate communication or motor signal i.e. an open loop BCI, or to be modulated adaptively in relation to a feedback signal i.e. a closed loop BCI (see Fig 1.1). Neural activity which could serve as a putative target for a communication signal might be an unrelated surrogate i.e. motor cortex activation extracted for a 'yes' signal, and visual cortex activation extracted for a 'no' signal (14). An alternative, more intuitive approach is to use activations linked to cognitive tasks which are more directly related to the process of communication. 'Cognitive' BCIs using neural activation related to the deployment of visuospatial and object based attention have been proposed (9,15), and are specifically investigated in this thesis. A specific example of how such a BCI might function practically is with a participant or patient using their neural activation to navigate a visual interface by deploying shifts of spatial attention, and then to select specific relevant visual exemplars such as a glass of water to indicate thirst, a particular face to indicate a desire to see a specific individual or to select a body part to indicate the need for attention or for a clinical assessment. I investigate the use of attention-driven BCIs, utilising visuospatial as well as object category-specific attention in Experiments 1, 2 and 3.

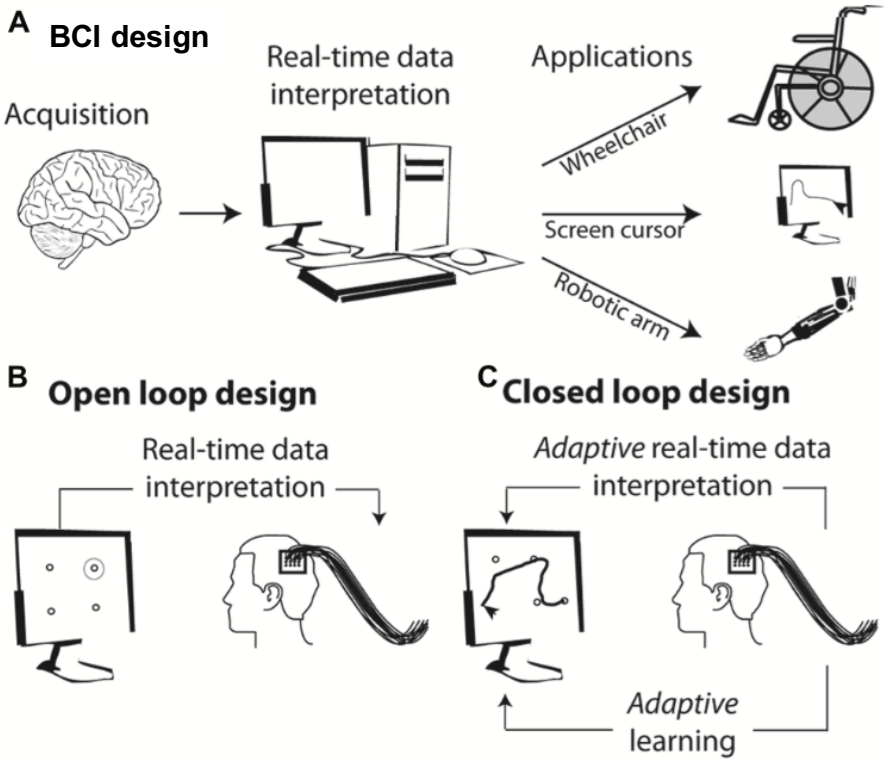

Fig 1-1. BCI schematic

A. Basic components of a BCI

B. Open loop BCI – brain activation is extracted form the participant, either invasively or non-invasively, without the participant being required to alter his/her neural activity.

C. Closed loop BCI – the participant actively engages with a visual interface which is linked to his/her extracted neural activity. The participant gains control of the visual interface by learning to control his neural activity in a target brain region/s.

Neural activation may be manipulated adaptively using neurofeedback, in order to train specific brain regions to modulate their functional output. This type of closed-loop BCI might be used to alter function in health e.g. increasing the ability to identify a particular face in relation to a detection task(16,17), or for the purpose of supporting or restoring function following a pathological insult e.g. engaging neural plasticity in a brain region affected by stroke(18), or a neurodegenerative disease

(13). In the final experiment of this thesis I examine the modulation of category-specific perception using neurofeedback training of higher order visual areas, and its effects on functional and structural plasticity.

As mentioned earlier a long-term goal of this research is to produce a BCI for therapeutic intervention, with the intended benefits of being practicable and portable. Rt-fMRI based BCI technology being non-invasive, can be used to optimise a specific brain- computer interaction, prior to conversion into a smaller, implantable invasive BCI for long term use. The use of rt-fMRI to delineate an anatomical substrate for direct implantation of intracortical BCIs has recently been demonstrated (10). Rt-fMRI 'verified' implantable BCIs e.g. (19) provide an optimal means of validating BCI technology prior to surgical implantation (20–23), increasing the safety and specificity of the BCI technique, as well as the likelihood of long-term success.

Target patient populations who may benefit from this type of BCI technology include those lacking the ability to communicate directly, either verbally and/or through the use of their limbs (24). Neuromuscular diseases such as amylotropic lateral sclerosis (ALS), a progressive disease of lower and upper motor neurones, ultimately lead to complete paralysis, whilst preserving cognitive capacity – a condition known as locked-in syndrome (25). All muscular control is lost, except for eye movement or minimal muscular activity, which is insufficient for independent communication. There are about 5000 people in the UK living with ALS, and it affects 1:100,000 people each year. Initially patients may use remaining muscular control to operate assistive communication devices, although as the disease progresses this become increasingly difficult. BCI technology has been successfully demonstrated as being useful in this patient group (26), as well as those developing LIS due to ischaemic brainstem strokes, with brain signals being used to control and enable the patient to communicate with their environment (27) .

1.1.2 fMRI as a tool for brain- computer interface

Following its introduction in 1990 (28,29), fMRI led to a revolution in cognitive neuroscience. It is based on the high structural resolution of MRI, with the 'functional'

element obtained from superimposed haemodynamic changes, thought to be produced by increased neural activity, linked to changes in brain function. Neuronal activity is tightly coupled with blood flow as a result of its dependence on oxygen. The magnetic susceptibility of oxygenated blood and deoxygenated blood differs, resulting in measurable differences in the magnetic signal produced when blood passes through a magnetic field. The governing principle of fMRI is the measurement of the Blood Oxygen Level Dependent (BOLD) signal of blood, a surrogate of neural activity – it is used when investigating cognitive processes, correlating it with processes related to concurrent performance of a specific task.

A challenge of conventional cognitive fMRI studies is that they only permit correlational inferences to be drawn between behaviour and BOLD activity. Stated more formally, in conventional fMRI studies, brain activity is the dependent variable, with behaviour being the independent variable. If this dependency were to be reversed, it might enable more direct links to made regarding behaviour and brain function. Seitz (30) states:

"A central goal of cognitive neuroscience is to understand how brains give rise to behaviour. The holy grail of many fields of cognitive neuroscience is to make causal links between the processing within, or between, various brain regions and people's perceptions, decisions or actions".

Taking this further, in the context of practical BCIs, fMRI might then allow brain states to be observed and manipulated to produce a specific behaviour.

Rt-fMRI, first described by Cox et al. (31), may be one way in which fMRI can be applied in order to allow for causal inferences. Rt-fMRI enables concurrent analysis and display of fMRI data, a process that is normally performed offline, over at least a period of days. The development of rt-fMRI arose from a number of concurrent advances in data processing including online motion correction (32,33), online multiple linear regression (34), cortex based analysis and visualisation (35), online analysis (36) and real-time paradigm control (37)

Aside from the attendant advantages of the technical innovations, rt-fMRI enables brain activation to be read 'on-the-fly'. The first opportunity that this provided was for

reading-out online brain activity in relation to specific cognitive actions i.e. an open loop BCI (see Fig 1-1B). This 'decoding' approach enabled a putative BCI based on the classification of brain activations to be conceived, and forms the first experimental thread of this thesis. The second opportunity that developed from the online analysis and representation of brain activation was its direct manipulation, with the aim of adaptively altering the level of brain activation (Fig 1-1C). This 'neurofeedback' loop can be used to train specific brain regions, with the aim of producing changes in brain function. It forms the second experimental thread of this thesis. Both of these approaches can be considered in the context of BCIs (17), and come under the stated aim of this thesis - to enable development of non-invasive communication and adaptive brain interfaces, for the translation of cognitive neuroscience techniques into the clinical arena.

1.1.3 Decoding

Decoding BCIs are conceptually the most direct application of a rt-fMRI based BCI. The character and pattern of brain activations have been shown to contain information relating to the function of the brain region being studied (38). Once a particular behaviour or cognitive process has been robustly linked with a particular brain activation, whether regional, or as a whole brain pattern of activity, the neural activations can be converted into bits of data which, for the purposes of a BCI, can serve as units for communication or information transfer (5). From here there is no requirement for an explicit behavioural output, as the imaging data can be used as a surrogate for communication. Using this principle, rt-fMRI allows haemodynamic brain activations linked to timed cognitive tasks, performed by a patient being scanned, to be interpreted 'on-the-fly' and used to power a communication BCI. A recent example of this was a spelling device, which linked brain activations produced by timed motor imagery, mental calculation and inner speech tasks performed by a patient, to the dynamic selection of letters from a virtual keyboard (39). Other practical applications of this method using rt-fMRI have included movement of robotic limbs (7), and computerised navigation devices (40).

For the purposes of this thesis, I investigated decoding of brain activations using rt-fMRI, as a means of communication in relation to the deployment of visuospatial attention. Attention is an ideal cognitive process for the implementing of a decoding BCI (15). Specific populations of neurons activate in response to the volitional direction of attention to circumscribed regions of space or features of objects in the real world i.e. endogenously directed spatial or feature-based attention. Objects and real-world stimuli at these locations also activate category-specific cortex, often in a 'spatiotopic' fashion (41).

The deployment of spatial attention occurs at the level of visual cortex, serving to enhance populations of neurones associated with retinotopically-represented regions of space in the outside world. As such, two potential sources of neuronal activity can be targeted for decoding in relation to driving an attention-based BCI. The first can come from representations of space in primary visual cortex, based on the underlying retinotopy, which are selectively acted upon by higher-order visual cortex. The neuronal populations in primary visual cortex are specifically enhanced, producing a signal which can be decoded and utilised (19,42–44). Andersson et al. (10) have previously shown that brain activations in primary visual cortex can be used to drive a rt-fMRI attention-based BCI. Participants maintained central eye fixation while deploying shifts of attention in four directions. The resulting brain activations were used to control the movement of a robot avatar (45). They further demonstrated that a non-invasive BCI could be used as an optimisation tool prior to implantation of a 'invasive' BCI, targeting the specific cortical location identified by the rt-fMRI paradigm for the implantation of an ECoG-based BCI (10).

An alternative approach could be to identify top-down 'command signals', produced in higher order visual cortex in relation to the control of internally driven processes e.g. endogenous attention. Candidate regions for the production of these signals include the parietal lobe, which has been suggested to have a 'salience map' (46,47). A cognitive BCI could potentially target such a region in the higher tiers of the visual hierarchy, in an attempt to provide surrogate control of this attention 'command signal'. In the experiments presented in this thesis, I find that regions with an intermediate position in the visual hierarchy such as the lateral occipital complex (LOC), for which there is also evidence of retinotopy and object-category information

representation, may provide higher rates of decoding accuracy for the direction of spatial attention, and serve as a potential target for a BCI-based intervention (48,49). An attention-driven BCI would be useful for patients who are unable to speak or move, but have preserved cortical brain function i.e. locked-in syndrome following brain stem strokes, or amyotropic lateral sclerosis. An implantable, intracortical BCI, using decoding of hand motor signals to control reach-and-grasp with a robotic arm, has been demonstrated in two tetraplegic patients with some success, although the long-term maintenance of these devices is an issue (22,50). Non-invasive BCIs have the benefit of being low-risk, with none of the attendant morbidity of surgical implantation, and could serve as a bridge to a portable system, by training and validating the use of specific brain activation decoding.

Following the identification of these brain activations for use in a BCI, the strength of the signal can be increased by training the BCI-user to volitionally modulate the level of brain activity. Andersson et al. (10) also used this in conjunction with the attention-driven BCI, encouraging patients to actively regulate the level of brain activation, while decoding was being performed. This second use of rt-fMRI as a 'closed loop' BCI system is called neurofeedback, which I examine in detail in the next section.

1.1.4 Neurofeedback

The other principal BCI application utilising a rt-fMRI framework is neurofeedback. Neurofeedback may be defined as the voluntary regulation of a neural signal in response to feedback of that signal, with a contingent reward for control of the signal. It was first established in the 1960s, building on work with biofeedback control of systemic autonomic measures, such as bidirectional control of heart rate (51). Control of a physiological measure led to the possibility of modulating the function of the effector organ, or in the case of neurofeedback, a target brain region, to produce long-term changes. For brain-based signals, early animal studies demonstrated that they could be controlled via neurofeedback to the animal. Fetz et al. (52,53) showed that for a reward, monkeys could selectively produce firing in neurones from which microelectrode recording were taking place and being fed back visually to the monkeys, i.e. without effecting whole brain function. Similarly Wyrwicka and Sterman

(54) showed that cats could self-regulate EEG activity from sensorimotor cortex. In humans, Kamiya showed that aspects of alpha band activity measured with EEG could be voluntarily regulated (55,56). However, these EEG-based BCIs for the purpose of neurofeedback, despite successful translation into the clinical environment, were limited by relatively poor spatial localisation of the electric source, even in the presence of modern multi-channel systems. FMRI-based neurofeedback, providing whole brain coverage and spatial localisation in the order of 1-2mm, was facilitated by the advent of realtime analysis software. Neurofeedback using rt-fMRI, is based on the online analysis of a participant's brain activation in the scanner, and concurrent presentation of a signal related to this activation back to the participant using a visual interface. This enables the voluntary training of one's own brain activity. Training a participant to voluntarily activate specific brain regions allowed neurofeedback studies to pursue altering a brain region's functional output, leading to associated cognitive and behavioural changes. Rota et al. (57) trained healthy participants to selectively increase the level of activation in the right inferior frontal gyrus, a brain region involved in the emotional processing of speech. Following training, participants were more accurate as compared with controls, in identifying the prosody of a speech task. No change was noted with regards to syntactic processing, confirming the specificity of the training effect. Other examples of neurofeedback success in healthy participants include modulation of sensitivity to pain with anterior cingulate and insular cortex control (58–61), memory performance after training of parahippocampal (PPA) activity (62), control of emotion and mood following modulation of amygdala (63) and anterior cingulate (ACC) activation (64), and motor performance with up-regulation of motor cortex activity (65).

The work in this thesis is intended to ultimately be used for translation into the clinical arena. Recent work with neurofeedback has shown promise with regards to its effectiveness in clinical application. An important study illustrating the strength of the technique in patients was conducted in five early-stage Parkinson's disease (PD) patients. Modulation of brain activation with neurofeedback training over a month, specifically of the supplementary motor area (SMA), led to clinical improvements in the patients' motor symptoms. Further, neurofeedback has been shown to result in enhanced network connectivity (66). This could be used to produce changes in functional brain regions related to but remote from pathology, enabling function by

engaging compensatory pathways. Disease processes involving cortical-subcortical networks such as PD are optimal candidates for this (13).

Other clinical conditions to which neurofeedback has been applied include chronic tinnitus control following down-regulation of auditory cortex (67), anterior insula control in criminal psychopathy (68), improvement of chronic pain measures with ACC modulation (59), and clinical improvement of unipolar depression following prefrontal cortex up-regulation (69).

As with the decoding approach with an attention-based BCI outlined above, the use of neurofeedback training can engage brain regions at various tiers of the cortical hierarchy. I specifically investigate the use of higher order visual brain regions for neurofeedback training in relation to biasing perception. Modulation of primary sensory cortex with neurofeedback has been shown to cause alterations in perceptual biases, although it was very tightly linked to the task used for neurofeedback training. For example, unconscious visual perceptual learning for specific orientations of contrast gratings was produced by exposing participants to the same grating orientations and then linking a feedback signal to the associated brain activity produced in primary visual cortex. Similarly, auditory discrimination was improved by neurofeedback training of participants on an auditory tone mismatch task (70). An open question, which I seek to answer, is whether perception can be biased using a more unconstrained approach, engaging higher order visual areas. In doing so, transfer gains produced by neurofeedback training may more generalised, and less limited. Additionally, I attempt to establish the role of higher order cortical regions in the control of neurofeedback training. This may provide further insights into the underlying mechanisms by which neurofeedback works (see next section). From a clinical perspective, neurofeedback-induced enhancement of higher order cortex may provide an alternative focus for therapeutic intervention. Neural recovery and rehabilitation are correlated with domain-general network activations in midline cortical structures, as compared to more localised changes. For example, speech recovery after stroke is linked to midline cortical structures (regions involved in the salience networks) (18,71), rather than more focal speech-related brain regions. These studies implicate the integral role of top-down ‘control’ regions, which may serve as a useful target in neurofeedback-directed enhancement of brain function.

1.1.5 What neural mechanisms underlie neurofeedback ?

The mechanism/s by which neurofeedback training results in learned self-regulation of brain activity is incompletely understood and is likely to be multifactorial, involving bottom-up and top down processes (see Fig. 1-1, and Table 1-1 below). The act of controlling a 'remote' feedback signal arising from contingent brain activation may engage multiple types of 'learning' and brain regions (18,72,73), and could be better described as a 'meta-learning process'.

An important unanswered question is whether there are brain regions which guide, or control neurofeedback learning. These putative brain regions would be commonly activated across multiple neurofeedback learning paradigms, and would interact with the specific brain region being modulated. Although some studies have examined the role of functional and effective connectivity (i.e. the directional coupling of brain regions determined a priori to be involved in a specific functional task) during neurofeedback training (e.g. 63–65) , very few have investigated the role of a possible control region or network which might co-ordinate the production of the neurofeedback training effect. Enriquez-Gueppert et al. using structural MRI and diffusion tensor imaging (DTI), demonstrated mid-cingulate volumes, as well as its white matter connectivity, predicted the ability of participants to control frontal-midline theta activity during EEG neurofeedback (77). The importance of white matter integrity, and therefore white matter structural connectivity was also illustrated by Halder et al., in relation to EEG based neurofeedback modulation of the sensorimotor rhythm using motor imagery (78). This has been supported by similar findings in behavioural training, with language acquisition (79) and video gaming skills acquisition (80). In Chapter 6, I explore the potential effective connectivity between trained regions and a putative control region, specifically the superior parietal lobe (SPL), using dynamic causal modelling (DCM). The role of SPL in the control of rt- fMRI neurofeedback training of visual areas has been recently suggested by Scharnowski et al. (81), who demonstrated improved visual detection of targets, following learned up-regulation of activity in primary visual cortex.

Most recently, a meta-analysis of individual participant data across 12 studies, 175 participants and 899 neurofeedback runs identified two specific regions, the basal ganglia and bilateral anterior insula, as being consistently activated (82). A subsample analysis in this study across 8 studies and 103 participants with an extended field of view, identified activation in a number of brain regions, including the dorsolateral prefrontal cortex, bilateral temporo-parietal areas including the SPL and lateral occipital areas. Deactivation was observed in the precuneus and the posterior cingulate cortex. Examining these activations within the context of known brain networks may help to shed light on the underlying mechanisms involved in neurofeedback. Some of the activated regions form part of the dorsolateral-parietal network also known as the central executive network (18,71), which is activated during demanding cognitive tasks, involving moment-to-moment monitoring of task performance, and manipulations of working memory and decision making. The anterior insular cortex and cingulate, together with the ventral striatum constitute part of the cingulo-opercular or salience network. This is a limbic- paralimbic system which deals with competitive, context-specific, stimulus selection, helping to focus the 'spotlight of attention' and enhance access to targeted neural resources required for goal-directed behavior (83). The third network which may be implicated here is the the fronto-parietal network, or 'dorsal attention network', and includes the dorsolateral prefrontal cortex, the frontal eye fields and the intraparietal sulcus. This network is responsible for maintaining a priority map of the visual environment, and may interact with higher order visual areas during processing of visual imagery and visual feedback.

Theories of the manner in which biofeedback control is developed have been put forward, and can be applied more specifically in the context of neurofeedback learning. 'Operant' or 'instrumental conditioning' (84) is frequently invoked as a key process in neurofeedback training (85,86). Learning occurs through a process of guided association, with a contingent reward for desirable behaviour, and/or punishment for undesirable behaviour. In rt-fMRI neurofeedback, the 'reward' can be provided by the feedback signal itself (e.g. through a visual interface), or by more explicit reward such as social affirmation (87) or monetary gain (88). In addition to the cortical regions used to provide the feedback signal, activation in the reward network has been reported. This includes anterior cingulate, orbital and dorsal

prefrontal cortex, and subcortical regions, including amygdala, caudate, and the putamen (i.e. (87,89).

Lacroix's 'two-step process' (90) suggests involvement of higher order brain regions and cognitive processes, with bottom-up control of a salient feedback signal during the active phase of learning. Repeated matching of the feedback signal with successful mental imagery strategies is mediated through interoceptive awareness. As this internal model improves in efficiency, it is maintained, and subsequently stored in implicit memory (72). In the final step the neurofeedback control becomes automatic (91).

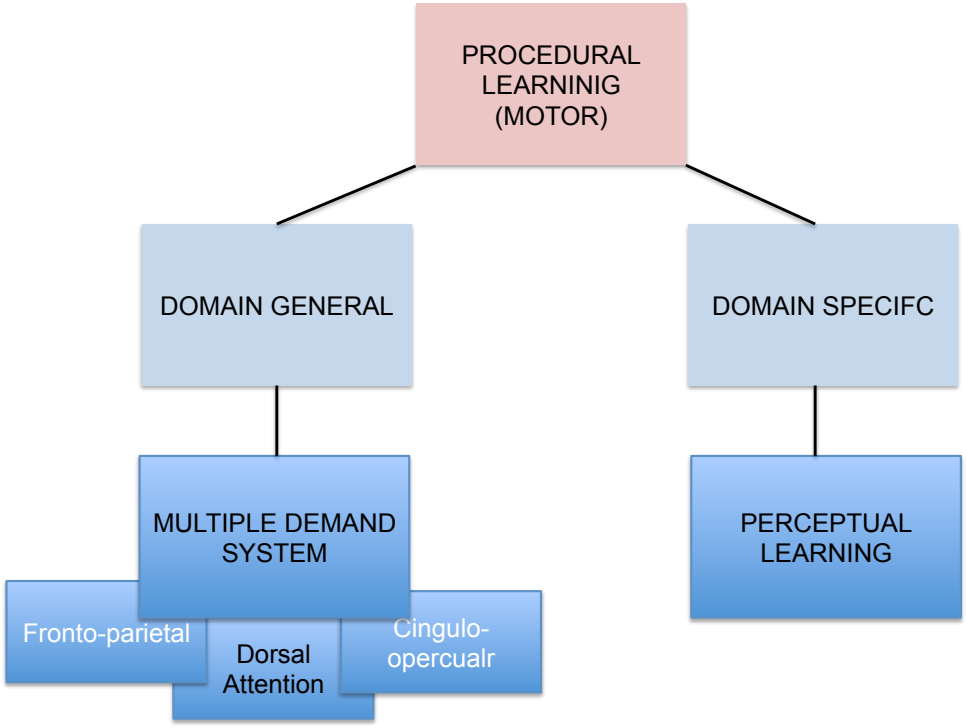

Fig 1-2. Types of learning engaged during neurofeedback

'Classical conditioning' is thought to be required for transfer of learnt regulation, and is likely to occur through an intrinsic or interoceptive cue that the participant learns to associate with the voluntary regulation of brain activity. In this regard neurofeedback may be a co-ordinated metacognitive processes, i.e. the participant learns by being aware of the specific associations between self-generated thought processes and control of the visual representation of the feedback signal. Garrison et al. (92) demonstrated differential levels of activation in posterior cingulate cortex (PCC)

during mind wandering and meditation; the PCC is regarded a central hub in the default node network, and is activated during mind wandering and self-referential processing. Using neurofeedback modulation, successful meditators were able to selectively down-regulate activation of the PCC. This study suggests evidence for the link between metacognitive, self-referential processes and learned modulation of brain activity using neurofeedback.

The involvement of corticostriatal regions as neurofeedback control regions has been suggested, supporting a 'motor theory' of neurofeedback (93). The original motor theory of learning is based on the observations of Lang and Twentyman, who taught participants to control their own heart rate using analogue biofeedback (51). The sequential or procedural element required (i.e. activity - movement – systematic, repetitive use of symbolic information) was conceptualised as being similar to learning a motor skill. Support for this theory in animals has recently been provided by Koralek et al. (94), by training rats to control a bi-directional auditory pitch in relation to two types of food reward. Food pellets were provided in response to a high pitch, and a sucrose drink was provided in response to a low pitch (see Fig. 1-2). The rats became proficient in both tasks. However, retaining the reward but removing feedback did not result in learning, nor did onset of satiety or removal of the food reward in the presence of accurate auditory feedback. Increased coupling between motor cortex and striatum in the 4-8Hz range, was associated with learning in the rats. Repeating the experiment with NMDA-receptor knockout rats, showed that despite intact overt movement, they were unable to learn voluntary self-regulation of brain activity. NMDA receptors are required for long-term potentiation in striatal neurons. These rats did not demonstrate increased firing in motor cortex and striatum with learning. This landmark experiment elegantly demonstrates the potential requirement of the motor system in the process of learning.

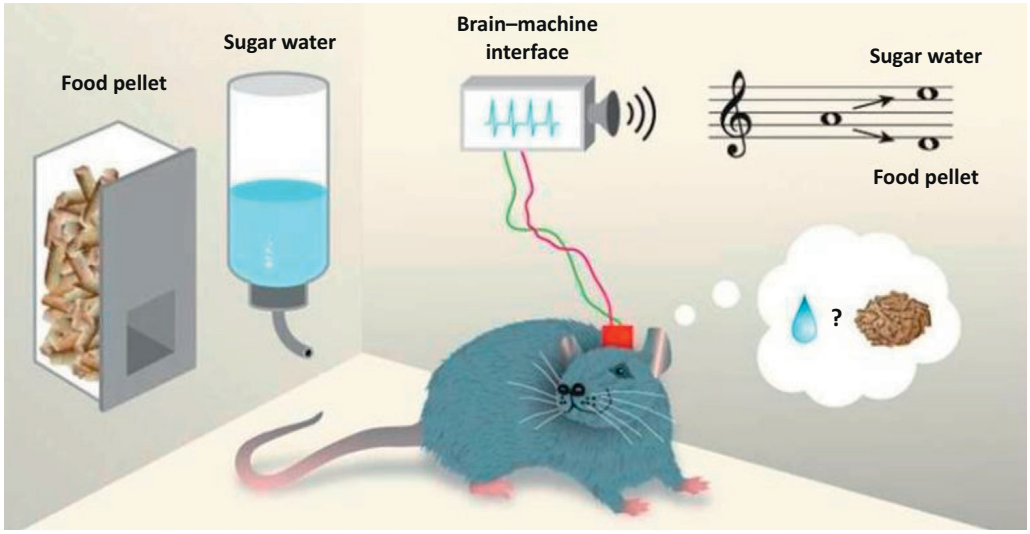

Fig 1-3. Implanted brain-computer interface

Rodents with electrodes implanted in their motor cortex, learned to control activity in relation to a bi-tonal pitch to secure either a liquid or a solid food reward (86,94)**.**

In summary, neurofeedback is a powerful but incompletely understood process. It is likely that the mechanism by which neurofeedback training produces neuroprosthetic control of a visual interface, engages multiple types of neural mechanisms, which differentially contributes at various stages of neurofeedback skill acquisition. Table 1-1, adapted from Strehl (72) relates each of the existing theories of neurofeedback control to each of the steps that are proposed to take place during neurofeedback training and learning.

In this thesis, I therefore set out to establish if the function of a specific brain region could be exposed in 'realtime' through first, decoding activity patterns from these areas, and then submitting them to the participant for training with neurofeedback.

Table 1-1. Mechanisms involved in skill acquisition during neurofeedback learning. (Adapted from Strehl (72)).

Paradigm	**Mechanism**	**Component of Neurofeedback training**
Operant conditioning	Trial and error	Reinforcer Shaping
Procedural learning	Motor repetition	Mental imagery Repetition
Perceptual learning	Percept repetition	Mental imagery Repetition
Learning outcome of a behaviour		
Classical conditioning	Target behaviour associated with conditioned stimulus	Transfer
Learning to predict contingency		
Motivation	Intrinsic +/- extrinsic	Individual differences
Two process theory	1/ Operant conditioning of adequate behaviour	Knowledge of result Instruction/ Strategy
	2/ Interoceptive association	Practice Feedback

1.2. Theme 2

Investigating higher order visual function

1.2.1. Introduction

A second theme in this thesis is the online investigation of the voluntary regulation of activity in visual brain areas to examine brain functions such as awareness and attention. Turning now to the mechanisms underlying human vision, the 'act' of vision and the epiphenomena that arise from it rests on three foundations (95):.

1. Optical information
2. Physiology and anatomy of the visual system
3. Phenomena of visual perception

The visual processing stream begins with retinal stimulation by photons emitted or reflected from the surface of objects in the environment. This 'optical information' is encoded and transferred along the optic tracts, to the superior colliculus and lateral geniculate nucleus, and then via the optic radiations, to the primary visual cortex at the occipital pole (95). Visual information can be categorised along a number of different dimensions (e.g. colour, space, form) and this is reflected by functional specialisations in a number of participating visual brain regions. Control of visual processing, together with the phenomenological aspects of vision, engage brain regions along a visual hierarchy, with higher order brain regions, such as the parietal lobe and frontal eye fields (FEF), acting upon incoming visual information being processed in the lower tiers, such as primary visual cortex. (See next section).

Two specific cognitive processes investigated in this thesis for the purposes of a BCI, targeting communication and neurofeedback respectively, are attention and perception. These processes are related, and frequently interact - the former provides a selective gating process for the brain-environment interaction, and the latter provides context and meaning to the information entering awareness/ consciousness. Higher order visual control processes are engaged in both of these

processes, although the manner in which they act upon lower order targets is incompletely understood, and has implications for the design of neuroprosthetic BCIs that seek to harness these processes. The construction of a rt-fMRI BCI allowed me to non-invasively investigate and manipulate higher order brain processes in relation to particular cortical regions in the visual hierarchy.

Before discussing these processes in more detail, I will expand upon the role of top down processes, and associated cortical regions in the context of the visual hierarchy.

1.2.2 Organisation of the visual hierarchy

Top-down processing in the brain works in the reverse order to the standard feed-forward sensory processing pathways. Progressively complex and anatomically distributed processing acts on simpler functions taking place in antecedent regions. It reflects a hierarchical organisation, and with some areas demonstrating functional subspecialisation (96).

Functional specialisation within the visual system includes colour, depth (V3a), motion sensitivity (areas V5/MT, MST, hMT+), and differentiation of object categories i.e. body parts (extrastriate body area, EBA), faces (occipital face area, OFA, FFA), places (PPA) (97–99).

More abstract and ‘gestalt’ visual processes are co-ordinated from ‘higher’ cortical locations in the hierarchy (100). Processes which require the involvement of higher order visual cortex, acting on lower and intermediate tiers of the visual hierarchy include attention, perception and prediction (96,101). The hierarchy of visual cortical areas begins in the primary visual cortex, and ‘feeds-forward’ information through two parallel circuits. The ventral circuit, including the inferior temporal cortex is called the ‘what’ pathway, and the dorsal circuit, which includes the parietal regions, is called the ‘where’ pathway (102). These two pathways are responsible for the processing of object and spatial-based information, respectively.

In a recent update to this two-stream model (103), Weiner and Grill-Spector (104) proposed a three-stream hierarchical model of the visual system (Figure 4) in which information processing pathways are extended beyond visual areas. According to this model, to convert vision to action, a dorsal occipito-parietal system connects the primary visual cortex with posterior parietal cortical regions. Regions implicated include the intraparietal sulcus (IPS), and are involved in position, motion, spatial working memory, form, attention and action. The traditional ventral occipito-temporal pathway is further divided into lateral and ventral pathways. The lateral stream passes through the lateral occipital sulcus/ middle occipital gyrus (dealing with position, form, shape), the posterior inferior temporal sulcus (dealing with position, motion), and the middle temporal gyrus (dealing with form, visual dynamics, touch, haptics, action and language). The ventral temporal cortex forms the first hub of the ventral occipito-temporal stream. This further subdivides, passing laterally to the occipital temporal sulcus (recognition of form, objects and faces) to the mid fusiform; medially is the mid-fusiform (convergence of perception to memory) to the parahippocampal gyrus. This model is therefore appealing because it attempts to 'complete' the processing pathways that start with the input of visual information and result in action, higher-order processing and/or storage in long-term memory.

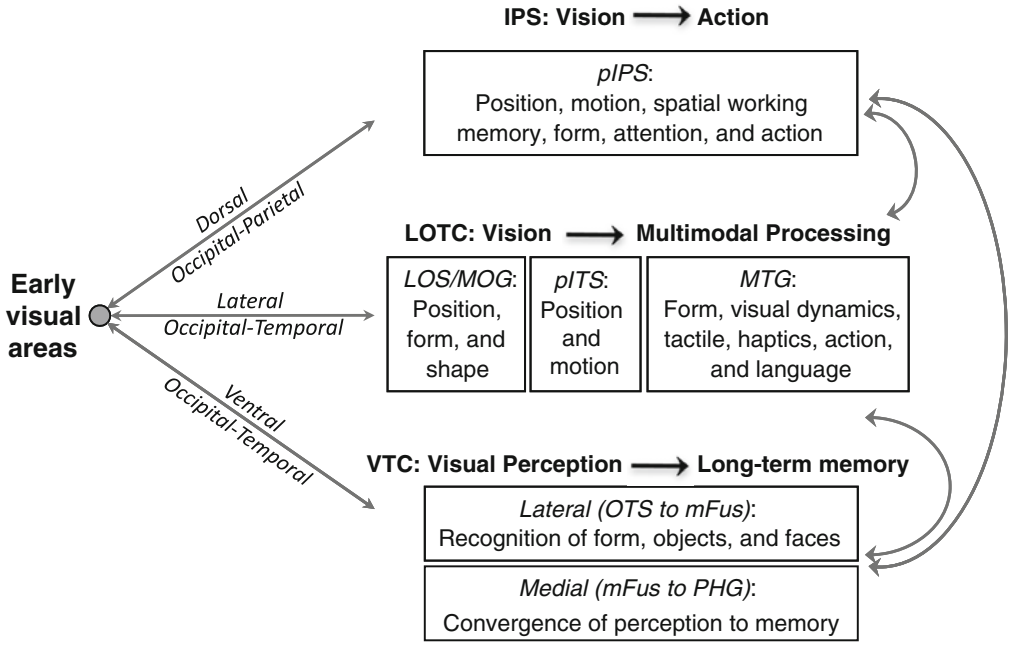

Fig 1-4. 3-stream model of the ventral hierarchy (proposed by Weiner and Grill-Spector (104)**)**

Intermediate regions in the visual hierarchy may serve to integrate aspects of the dorsal and ventral visual streams. As such these regions are likely candidates to host a salience or priority map of environmental targets, and importantly for consideration of a BCI, are ideal regions for extracting information. The parietal lobe for example has been shown to contain both spatial and feature-based information; anterior IPS (105) and the superior parietal lobe (106) have been implicated in the attentional binding processes (107–109) and can be activated during spontaneous and automatic 'gestalt' perception (110), a process which has been suggested to occur prior to attentional selection (111). Lateral Occipital Cortex (LOC), and ventrally placed object category selective cortex such the FFA and PPA, have been shown by some studies to demonstrate retinotopy (41,49,112). Further, these three regions preferentially respond to specific object categories but will demonstrate a range of sensitivities, albeit with a reducing gradient of activity. That is, FFA does respond weakly to non-face stimuli compared to face stimuli, and similarly PPA also responds to non-scene/house stimuli more weakly than scenes and houses. The presence of spatial as well as category-based information in regions such as LOC and FFA (113) may indicate a more significant role in more complex perceptual actions (100, 91). Having established a possible functional structure of the visual hierarchy, I will next consider the manner in which the two processes examined in this thesis, attention and perception, interact with this functional assembly.

1.2.3 Attention and top-down control

Attention is cognitive process that enables directed processing of specific aspects of sensory signals evoked by environmental stimuli. It may selectively involve any sensory modality, and may serve to 'bind' component information. The dominant view of how attention works mechanistically is summarised by Noudoust et al. (116):

"*The objective of attention can be viewed as increasing the signal-to-noise ratio (SNR) of the readout from sub-populations of neurons encoding the selected representation"*

This preferential processing can be facilitated through bottom-up, or salience driven processes and top-down, or stimulus-driven processes. Top down attention can be separated into spatial attention, which utilises cortical retinotopic topographical representations of the external world, and feature/ object based attention that enhances visual features of an object, or whole objects. Although specific cortical circuits sub-serve the different aspects of top-down attentional control (117–120), there is a significant degree of overlap (48,112). It has been suggested that one or more high order regions may generate an 'attentional command signal' which biases spatial and non-spatial features, as well as integrating emotional and motivational valencies, by acting on a attentional priority map. Such a map would contain top-down weighted salience-based representations (101).

Candidate cortical regions for generating a 'command' signal for the allocation of visuospatial attention include frontal and parietal regions, although work in this area is still very much on-going. FEF has been shown to have a causal role in both spatial and feature based attention (121). Microstimulation and lesion studies in monkeys (122), and fMRI studies in humans (105,118,123) have identified the FEF as a potential source of an 'ultimate attention command signal'. Parietal regions implicated include the intraparietal sulcus (117), shown in studies on normal participants (124), lesion patients (125–127) and with TMS (128–130). The directionality of the frontoparietal relationship is contentious, with evidence for IPS directing information to the FEF (131) and vice versa (129,132,133).

Identifying a 'control' region for visuospatial attention is not only of neurobiological significance, but is advantageous in the implementation of an attention-based decoding BCI (15). Decoding brain activations from such a region in relation to the deployment of attention may be more likely to yield successful results, as a result of a higher signal-to-noise ratio.

1.2.4 Visual Perception

Perception encompasses a number of interlinked processes that produce the contents of consciousness, and give rise to visual awareness (134–136). The act of processing visual information from the environment is influenced by expectation, and previous exposure (e.g. priming, training) (30,137–140). The manner in which these processes alter perception, by acting on perceptual brain regions prior to the subsequent act of perception can be linked by the changes in the level of brain activations (141). This is particularly relevant to the neurofeedback study in this thesis (Chapter 6), which aimed to alter perception by modulating the levels of activation in brain regions involved in visual perception. Electrophysiological and neuroimaging techniques have been used to quantify the levels of neural activity and its link to the mechanisms underlying perceptual phenomena (142). Single neuron and local neuronal population activity, as measured with increases or decreases in spike rates, have provided insights into perception (143), but are constrained by being limited to very small cortical locations. Perceptual phenomena are likely to result from dynamic, hierarchical interactions between multiple cortical and subcortical regions; fMRI, by measuring across the whole brain, has been useful in demonstrating changes in brain activation in primary retinotopic cortex, category-specific brain regions, as well as high order regions, in relation to visual perception (144,145).

Increased brain activity in primary visual cortex, linked to the occurrence of specific percepts, results in improved visual discrimination (17). The effect of prior exposure on subsequent perception may serve to preferentially enhance a specific perceptual state. For example, previous exposure to similar stimuli via priming with degraded objects results in increased functional activation in visual areas as compared to priming with intact stimuli; the latter results in repetition-related decreases in BOLD signal (146,147). Mental imagery of stimuli produces activation in category-specific areas similar to the physical presentation of the stimulus, although the level of activation is quantitatively lower (148,149).

Activation in areas such as category-specific regions and the involvement of top-down control areas may facilitate category-specific changes in perceptual state – a

concept fundamental to the study presented in Chapter 6. During bistable perception of faces and houses, prestimulus activity in FFA but not PPA predicted which of the two percepts would subsequently be perceived (150,151). A top-down influence over object identification was demonstrated by Eger et al. (152) who showed that the threshold for the identification of gradually-revealed visual objects was lowered by providing verbal cues, and was associated with an earlier rise of mean BOLD activity in ventral visual areas, as well lateral parietal and prefrontal cortex. The authors concluded that recognition was facilitated by fronto-parietal coupling with ventral visual areas, indicating top-down control of perception. Effective connectivity analyses, using both Dynamic Causal Modelling and Granger causality, revealed directional BOLD signal flow from parietal to occipital cortices that is, feedback during a visual imagery task. Flow reversal, that is, feed-forward signalling, instead occurred during a viewing/ simple perception task (153). Previous work by Mechelli et al. (154) demonstrated segregation of top-down effective connectivity in relation to different perceptual states. Superior parietal cortex demonstrated increased effective connectivity during both imagery and perceptual tasks, whilst prefrontal cortex was implicated only during visual imagery, with specificity for category-specific cortex. These findings suggest that more demanding perceptual phenomena increase top-down involvement. In keeping with this, bistable perception involving ambiguous stimuli engages top-down processes more than simple, unambiguous perception (155).

Attempting to successfully modulate perception, as I set out to do in Chapter 6, benefits from a conceptual understanding of the mechanisms that underlie it. Helmholtz proposed an inferential process, with the requirement for ‘unconscious inference’ to inform visual input through the eyes, in order for the brain to make sense of the visual scene (156,157). The idea was successively refined (158,159), and subsequently set out within a Bayesian hierarchical model (160–169), where prior knowledge of the environment is stored in the brain and modulates incoming sensory input to produce the perceptual experience. The level of brain activation in neurons sensitive to expected knowledge has been associated with short term learning of expectations and perceptual ‘priors,’ with findings of increased efficiency/ decreased activations in fMRI studies, and higher levels of activity in neurophysiological studies. Taken together, the level of activity within a brain region

may represent prior distributions of neural activity, which have been shaped by sensory input and expected sensory inputs; modulation of the level of this activity using for example neurofeedback, may therefore impact on the subsequent processing by that region, and its resulting output. Therefore, considering perception as an inferential process, one approach to producing perceptual changes would be to alter the prior state of the brain regions involved in perception. This is supported by evidence demonstrating the influence of expectations on perception. The interaction of neurofeedback and visual perception in Chapter 6 provides a direct means of doing this, as well as providing a novel insight into the relationship between the brain and visual perception.

1.3 Summary of experimental aims

The thesis sets out to develop a non-invasive BCI, powered by cognitive processes, which target higher order visual cortex. Over the course of the experiments in this thesis, I utilise novel technical innovations to achieve this and gain insights into the neurobiology of the underlying neural processes.

The first aim of this thesis was to decode visual information in higher order visual cortex to identify the focus of an individual's spatial attention. In Experiments 1 to 3, I establish offline, and then online methods for achieving this, resulting in the implementation of an attention-based BCI.

The second aim of this thesis was to implement rt-fMRI for the purposes of a neurofeedback BCI, in order to bias visual perception by training participants to modulate category-specific visual brain regions. In Experiment 4, I test this aim using a binocular rivalry paradigm with face and house stimuli, employed before and after participants learned to modulate a neurofeedback signal based on the activation levels in FFA, and PPA. Face and house stimuli robustly activate FFA and PPA respectively, and these brain regions were selected on account of their well -

documented role in face and house processing, both during the deployment of object based attention (99,170) and in category-specific perception (151,171).

In anticipation of the experimental findings, I demonstrate that higher order visual areas can be successfully used as the basis for a rt-fMRI BCI. I was able to successfully decode covert deployment of spatial attention across four quadrants on a single trial basis using a novel automated decision criterion. I also show that participants' perception of faces or houses was biased following rt-fMRI neurofeedback training of PPA and FFA. The perceptual biasing was consistent with the specificity of the each participant's neurofeedback training signal. In sum, the results from these experiments inform the construction of effective patient BCIs.

2. GENERAL METHODS

2.1 MR PHYSICS

2.1.1 Introduction

Magnetic Resonance Imaging (MRI) was first independently demonstrated in 1946 by Bloch (172) and Purcell et al.(173). This work was based on the observation that atomic nuclei (of which all matter is composed) have magnetic properties, produced by a charged particle, namely a proton, spinning around is own axis (174). Following on from this, pioneering work by Lauterbur (175), Mansfield (176), and Damadian (177) enabled the formation of two-dimensional images, by identifying MR signal location, and heralding the application of topographical methods to MRI imaging in humans. The first image took 5 hours to perform in a human.

2.1.2 Superconducting magnets and coils

MRI scanners used today rely on the production of a magnetic field by running an electric current through a super cooled (i.e. -269 C) wire. The coolant used is cryogenic liquid helium. The wire is housed around a large diameter open bore, together with a number of other components related to maintaining the appropriate temperature conditions.

Another important hardware component of the MRI scanner are the *radiofrequency coils* (RF), which transmit and receive the radio-frequency waves used in image acquisition. There are a variety of coils e.g. volume coils, surface coils, which are generally shaped to ensure uniformity of the RF field inside the coil. The coils maybe receive only, or as in the case of a head coil, transmit-and-receive coils.

The aim of an MRI scanner is to produce a constant and relatively homogenous magnetic field inside the bore. The field is termed B_0 and it provides optimal

conditions for the excitation of protons, found in all organic matter likely to be placed in the scanner including humans, and allows optimal measurement of the energy transactions which occur as a consequence (178).

The RF field is referred to as the B_1 field. It combines with the B_0 field to generate MR signals that are spatially localised, and further encoded by the gradient magnetic fields to create an MRI.

The homogeneity of the scanner's B_0 field is important specifically for field localisation when scanning participants. Placing an object into the B_0 field however creates local perturbations/susceptibility effects. To counter this, *shim coils* provide field correction through active shimming (cf. passive shimming which is performed during installation using sheets of metal). Shimming therefore refers to changes made to the magnet in order to optimise field homogeneity.

2.1.3 The magnetic field, protons and 'wobbling'

The magnetic field to which the participant or patient is exposed to inside the scanner is measured in Tesla (T) – one T is approximately 20,000 times the strength of earth's field. All experiments in the thesis were performed in a 3T field (Allegra Siemens 3T magnet).

The source of the MR signal in the majority of human imaging studies is from the hydrogen atom, or more specifically a hydrogen nucleus. This consists of a single positively charged particle, a proton 'spinning' around the nucleus. This 'spin' is an intrinsic quantum mechanical property of elementary particles, and is an angular momentum, rather than a simple rotation. It is determined by two quantum numbers, one for magnitude (nuclear spin quantum number) and the other for the direction in the z-axis (azimuthal quantum number). Protons essentially behave as miniature superconductors in which the spinning positive charge produces an electric current, which in turn generates a miniature magnetic field. When atomic nuclei are placed in the scanner's magnetic field, the magnetic moments of the nuclei, which are

normally random, tend to align parallel or antiparallel to the scanner field, B_0, the axis of which is parallel to the bore of the magnet (z-axis).

The optimal state is at the lowest potential energy i.e. parallel. However, when placed in the B_0 field, the group of protons, regardless of their individual alignment, 'wobble' in a particular way (the classic analogy is a spinning top which wobbles but does not fall). This is called 'precession', and the rate at which this occurs is termed the precession frequency or 'Larmour' frequency, for a particular magnetic field, B_0. It is dependant on B_0, and this relationship is realised in the Larmour equation:-

$$\omega_0 = \gamma B_0$$

ω_0 is the angular frequency of precession of protons in an external magnetic field, γ is a constant for a specific nuclear species (e.g. hydrogen), and is termed the gyromagnetic ratio. Its value for a proton is 42.6 MHz/T. Therefore, the Larmor frequency f_0 for a 3 T magnet is 3(T)*42.56(MHz T^{-1}) = 127.68 MHz.

2.1.4 Longitudinal magnetisation

Protons 'precessing' parallel to B_0 gradually come into alignment to create a sum longitudinal magnetisation i.e. along the z-axis of the external magnetic field, which is termed M_0. Given that the participant contains the bulk of these protons, they essentially become a 'magnet'. The stronger the magnet, the more protons will align in the direction of B_0. However, in order to measure these specific protons, magnetisation, which lies at an angle to B_0, must be produced.

2.1.5 RF pulses and transverse magnetisation

An RF pulse applied across the B_0 field (and the participant in the scanner) knocks the protons out of alignment with the z-axis. This is due to a transfer of energy from

the RF pulse, and can only occur at the precession frequency of the protons – this provides for the 'resonance' term in MRI.

The RF pulses in fact produce two changes in protons, which give way to the images typically used i.e. T1 and T2 images. The transfer of energy results in some protons gaining a higher energy state ('up' state), being antiparallel to B_0 – this will result in net cancellation of the longitudinal magnetisation in some regions. Correspondingly, some protons will move *in phase* - this produces a transverse magnetisation i.e. x-y plane. This transverse magnetic field rotates at the Larmour frequency, inducing an alternating voltage in a neighbouring conductive (RF) receiver coil. This generates an electric current, which leads to the MR signal.

An important effect is produced by now switching the RF pulse off – the excited, 'wobbling' protons, *relax*, fall out of phase, and return to a lower energy state.

2.1.6 T1 and T2 relaxation

T1 relaxation occurs when the longitudinal magnetisation reduces, with the T denoting the time taken for the protons to recover thermal equilibrium. The T1 value is tissue specific, and is longer at higher field strengths. The value is dependent on the 'tumbling' rate or degree of molecular motion of different molecules. Free water, being unbound, and of a small molecular size has a very rapid tumble rate, indeed too fast to have efficient T1 relaxation. Large macromolecules e.g. membrane lipids, which have highly bound hydrogen, tumble slowly. Both free water and bound hydrogen have long T1 relaxation times. On other hand partially bound water, or substances such as fat have short T1 values i.e. close to the Larmour frequency (179,180).

A T1 curve describes the continuous and exponential process of the return of a proton to its original energy state – 'T1' is the time taken for longitudinal magnetisation to return from 0 to $(1\text{-}e^{-1})$ or 63% of its final value.

T2 relaxation occurs when transverse magnetisation reduces. It is once again tissue-specific, but unlike T1 relaxation, is unaffected by field strength, and may proceed with or without overall energy loss.

Two processes cause loss of transverse phase coherence in the x-y plane. One is inhomogeneity in local magnetic fields within tissue, which produces loss of phase coherence within a group of protons, so-called *spin-spin relaxation*.

There is also inhomogeneity within B_0, resulting from slightly different Larmour frequencies at different locations within the field. The dephasing that is caused by these field imperfections is a constant phenomenon, and is potentially reversible. T2 'star' (T2*) relaxation describes the combination of these effects, and determines the actual rate of decay of a signal, which is not affected by a gradients i.e. free induction decay (FID). FID is occurs in relation to decay by a number of processes including the magnetic field gradients used to localise and encode the MR signal. It is more usual to generate the MR signal using the application of a RF pulse, and measuring the signal produced in response by the spin of the protons – termed the 'echo'.

T2 is the time taken for transverse magnetisation to decay to e^{-1}, 37% of its initial value. The major determinant of this value is spin-spin interactions. Free water, which has its molecules relatively far apart, will have fewer of these interactions- it has a longer T2 value compared to tissue containing water e.g. the brain. Transverse relaxation occurs rapidly in human tissue and therefore T2 values are less than or equal to T1.

2.1.7 T1 and T2 weighted images

T1 images are produced when signal intensities are governed principally by differences in tissue T1 relaxation time – the manner in which this is achieved is by changing the time between two RF excitation pulses, the *repetition time* (TR).

T2 images typically have long TRs and are influenced by the differences in proton density and T2 relaxation times.

After the initial 90 degree RF pulse, the application of a further 180-degree RF pulse rotates the protons through 180 degrees – they precess in the same plane but the opposite direction. The signal amplitude of the MR signal is determined by the number of protons that are in phase following an application of an RF pulse- the maximum amplitude is reached at the *echo time* (TE), where TE is the time interval between the initial 90 degree RF pulse and the echo. In order to achieve the maximum signal the 180 degree RF pulse must be applied at time TE/2.

The signal produced following the application of the 180-degree RF pulse is termed the spin echo (SE) – that is the signal caused by the echo or spin of the protons in relation to the RF pulse. Persisting inhomogeneities in B_0 cause protons to continue losing phase coherence following a 180-degree RF pulse; repeated RF pulses will produce further SE's, which reduce in amplitude due to the T2 effects described. The T2 curve is formed by connecting the SE intensities as compared to the T2* curve, which shows the de-phasing effects in the absence of the 180 degree RF refocusing pulses.

Alterations in TE and T2 weighting alter signal contrast, but can also potentially alter the signal-to-noise ratio.

2.1.8 Contrast

The term contrast in this context applies to the process of differentiating adjacent tissues. It is determined by signal intensities, which relate to a number of variables including blood flow, the pulse sequence, TE, inversion time, TR, proton density and to tissue T1 and T2 relaxation times.

The relationship between these variables and signal intensities is complex. For example a much stronger signal is achieved with shorter TEs. However, when this is

less than 32 milliseconds, differences in T2 have less influence on tissue contrast. T2 images, as produced for cognitive neuroimaging are obtained using long TEs. Heavily weighted T2 images are useful for delineating tissue/water interfaces in diagnostic neuroimaging e.g. to image the trigeminal nerve.

Contrast media can also be used, and differences in magnetic susceptibility can be taken advantage of e.g. blood products, and iron content (181).

2.1.9 SE and GRE

The emergent concept is that the RF signal is not measured directly, rather MR imaging exploits the 'echo' that is produced in relation to it. I have discussed the formation of the spin echo and the relevant related measures above.

SE sequences have relatively long TR and TE values, and as a result are time consuming. Gradient echo (GRE) sequences provide an alternate approach. A single RF pulse is used, with smaller flip angles (5-40 degrees) compared to the 90-degree RF pulse used in SE. The 180 refocusing pulses are replaced with magnetic field gradients, which change field strength and alter the Larmour frequency in a specific direction. Protons in this direction rapidly fall out of phase, causing a sharp decline in the FID signal. The gradient 'echo' is produced by applying a second magnetic field of the same amplitude to the opposite direction, pushing the 'dephased' protons back into phase.

The combination of short TE and a RF with flip angles <90 degrees means that the preceding longitudinal magnetisation is not entirely abolished, enabling good signal quality with short TRs. GRE sequences are however quite vulnerable to magnetic susceptibility effects, with greater signal loss compared to SE sequences.

A final important point with GRE imaging as compared to SE, is that the gradient reversals (as opposed to the 180 refocusing pulses) do not minimise local field inhomgeneities – it only acts to refocus those protons that have been dephased by

the gradient itself. The result is that GRE sequences are more influenced by T2* effects, as compared to T2 relaxation only (181).

2.1.10 Image construction – localisation and encoding

A three-dimensional MR image is produced using 3 gradient fields (181). The position of the protons in the object being imaged are spatially localised in three dimensions using these fields, allowing an MR image to be constructed.

Slice-selection gradient

A single slice is localised by altering B_0 along a chosen axis using a gradient field; the Larmour frequencies of the protons varies within this gradient field. Manipulating the 'steepness' of the gradient field, or the bandwidth of the RF pulse can change slice thickness. Usually the RF pulse is applied as range of bandwidths, which excites a specific slice thickness. The following two gradient fields allow pixels within a 'slice' to be spatially encoded.

Phase-encoding gradient

Protons may be differentiated according to their phase following the application of a gradient magnetic field and despite having the same precession frequency, they will no longer be in phase. The gradient field, G_p is therefore applied in the 'phase-encoding' direction; the resultant spin phases will vary linearly across the phase-encoded direction.

Frequency-encoding gradient

In a similar fashion, the 'frequency-encoding' gradient is applied orthogonal to the phase-encoding gradient. Protons may now be differentiated according to their rotation frequencies in this plane.

To summarise, the slice selection gradient produces the slices that comprise the brain volume, and the phase encoding and frequency encoding gradients are

orthogonally applied to spatially document the position of the protons in each slice volume, to produce the final 3D MR image.

To ensure sampling signal i.e. the maximum 'echo', and to counteract de-phasing produced by transverse magnetisation, additional gradient pulses are used, just before the slice-selection gradient and immediately before the frequency-encoding gradient.

2.1.11 K-space

This term relates to the data matrix produced through successive and incremental applications of the 'pulse sequence'. It is composed of image localising gradients. Each pulse generates a single echo signal, enabling all of the pixels within a slice to be systematically localised; at each step, although the phase-encoding gradient is increased, the slice selection and frequency-encoding gradients are maintained. It is this data-matrix of spatial frequencies that is converted into an image using a Fourier transform (179).

2.1.12 Image acquisition

The total image acquisition time is the product of the time interval between pulse sequences (TR) and the number of phase encoding steps (N_p). Conventional pulse sequences e.g. SE and GRE, acquire one phase encoding step per TR, which is a single line of k-space. Faster image acquisition can be implemented, with multiple lines of k-space per TR being acquired e.g. Turbo SE or Echo-planar imaging (EPI) (181).

2.1.13 Echo-planar imaging (EPI)

EPI is one such rapid sequence, with the generation of multiple GREs per TR, through high amplitude gradients and rapid oscillation. Single-shot EPI acquires all of

the requisite phase-encoding steps in a single TR, whereas multi-shot EPI uses a few TRs. EPI sequences are useful for functional neuroimaging, as the T2* contrast is sensitive to blood-oxygen level dependent signal (BOLD) – all of the functional imaging in this thesis was performed with EPI sequences.

2.1.14 Magnetic susceptibility

The term denotes the extent to which a substance becomes magnetised when placed in an external magnetic field. Electrons within a substance, specifically those that are delocalised, or in nuclear orbit, interact with the magnetic field to produce circulating currents. These currents in turn generate local internal magnetisation, which either augments or opposes the external magnetic field. This results in either positive susceptibility in the first instance, or negative susceptibility in which the local magnetic field is reduced. Materials with positive susceptibility are termed paramagnetic, whereas those with negative susceptibility are termed diamagnetic.

Living tissue and water are very weakly diamagnetic, as compared to iron and other metals, which are very susceptible and are called ferromagnetic materials. Superparamagnetism describes an intermediate state of positive susceptibility that lies between paramagnetism and ferromagnetism.

2.2 The BOLD signal

2.2.1 Applying MR physics to produce a functional signal

Susceptibility, as introduced above, is exploited to deliver the functional component of neuroimaging, utilising the molecular relationship of oxygen to haemoglobin, the oxygen-carrying component of blood. This gives rise to the BOLD signal (182,183), which is the basis of fMRI.

Oxygen, a primary requirement for aerobic metabolism, is reversibly bonded to an iron atom at the centre of the haem component of haemoglobin. When haemoglobin arrives in tissue from the arterial circulation, oxygen separates from the haem molecule, exposing electrons from the iron atom. These unpaired electrons alter the susceptibility of the deoxygenated haemoglobin molecule, making it paramagnetic (28). Deoxygenated haemoglobin increases the spin phase dispersion, thereby increasing local field inhomogeneities. T2 and T2* relaxation is reduced, resulting in a decreased T2* weighted signal, and a reduced BOLD signal. Oxyhaemoglobin on the other hand, is diamagnetic, and consequently does not interact significantly with the magnetic field (184). Note the described susceptibility effects relate to intravascular changes. Although susceptibility effects will occur inside blood vessels and adjacent tissue, at field strengths of 1.5 T and above, more than half the BOLD contrast is due to the intravascular signal (185).

Briefly considering the extravascular component, the largest BOLD changes (i.e. changes in T2*) occur near veins in the region of an ‘activated’ area – this is related to the comparatively reduced water diffusion around veins, as compared to capillaries. With diffusion, water molecules move randomly, sampling a range of spatially differentiated magnetic fields. The net phase dispersion between different spins is reduced, leading to a comparatively smaller change in T2* relaxation. Capillaries allow for rapid water diffusion, and so BOLD changes are less pronounced (186,187). Note, however that despite this, capillary dilatation accounts for 85% of the increase in blood flow linked to neuronal activity, and provides the

trigger for arteriolar dilatation (see next section). It becomes apparent that while the term BOLD implicates oxygen, it is better regarded as a measure of deoxyhaemoglobin washout, and the magnetic susceptibility changes that accompany it.

Aerobic metabolism, which is closely linked to neuronal activation, should result in a rapid fall in the level of oxygenated haemoglobin through oxygen extraction for cellular metabolism – the assumption that follows is that increased neural activation would cause a decrease in the T2* derived MR signal. However the converse is observed i.e. increased BOLD signal/ brighter T2* weighted images (28,188). This is because there is a disproportionate delivery of oxygenated blood as compared to that required for cellular metabolism, in response to increased neural activation (189).

2.2.2 Physiological basis of BOLD

The link between neural activity and the measured changes in deoxygenated haemoglobin is complex and is underpinned by the following:

- the type of neural activity and the cells generating it
- the link between activity and energy demands
- the processes coupling demand, supply and energy metabolism

I examine these steps in detail below.

2.2.2.1 The haemodynamic response and the haemodynamic response function

The time course of the BOLD response is termed the haemodynamic response. This can be modelled using the haemodynamic response function (HRF). The function, which in SPM is a linear combination of two gamma functions, is convolved with a stimulus function to produce a task-related regressor (see chapter 2.4 for how this is achieved in general linear modelling.) In terms of the underlying physiology which

gives rise to the haemodynamic response, it is thought to arise in the following way: in response to a task, neural activity occurs very rapidly, on the order of milliseconds. The associated vascular changes occur more slowly, in three stages listed below, giving rise to a HRF that peaks 4-6 seconds after the onset of neuronal activity (190).

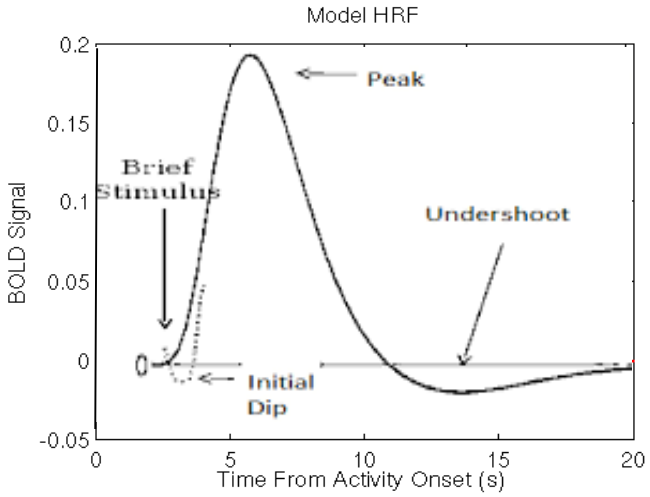

Fig. 2-1: Illustration of the HRF. Each of the stages described below are labelled.

1) **INTIAL DECREASE IN HRF**

   (due to increased oxygen consumption i.e. the primary inflow and volume effects)

Recent, work (191) suggests that an initial decrease is produced by capillary dilatation, which precedes arteriolar dilatation, in response to neural activity. Transient reduction in blood volume is produced by the very rapid capillary dilatation, which is mediated by capillary pericytes actively relaxing in response to neuronal activity

Following this, on-going capillary dilatation may generate as much as 84% of the subsequent steady-state increase in blood flow evoked by neuronal activity.

2) **LARGE INCREASE IN HRF above the baseline**

This has been explained by capillary dilatation, and the slightly later arteriolar dilatation, which is then accommodated by venous and venular vasoldilatation – the

latter, due to their ‘balloon-like’ elasticity, results in the housing of a large volume of predominantly deoxygenated blood.

3) **DECREASE IN HRF with an UNDERSHOOT below baseline i.e. ‘post-stimulus undershoot**

The undershoot lasts for several seconds, and can go on for as much as half of the haemodynamic response. However, it is inconsistently observed in the average response, due to factors such as variability in its shape and amplitude, which can vary on a trial-by-trial basis.

The origins of the undershoot are incompletely understood (192). It has been suggested to reflect the metabolic rate of the tissues, measured in terms of oxygen consumption ($CMRO_2$), with neuronal activity producing a sustained increase in $CMRO_2$. The ‘metabolic’ theory suggests that the raised $CMRO_2$ then causes a reduction in the levels of oxygenated haemoglobin, resulting in a drop in the BOLD signal. A ‘compliance’ theory suggests that the dip is due to a persistence of increased blood volume following the cessation of brain activity. It is likely that there is a contribution of delayed vascular compliance (192).

Interestingly, both post-stimulus ‘overshoots’ and ‘undershoots’ have also been recorded. Current work indicates the post-stimulus undershoot is accompanied by cortical inhibition, and reductions in cerebral blood flow (CBF) and $CMRO_2$, with the opposite being true in the event of a post-stimulus overshoot (193).

2.2.2.2 Energy coupled blood flow, and blood volume changes

As discussed, the BOLD signal, in response to a stimulus (and the resulting neural activity), is produced by a decrease in deoxygenated haemoglobin, reflecting regional oxygenation. In order for this to take place, commensurate changes occur in CBF, cerebral blood volume (CBV) and cerebral oxygen metabolism $CMRO_2$, itself a surrogate measure of glucose metabolism (194). The manner in which these measures are effected is related in part to the relative contributions of the neural

activity, which include spiking and synaptic activity, and may be excitatory or inhibitory activity (195). The interaction of aerobic and anaerobic metabolism in neurones and their neighbouring cells (e.g. glia), impacts on this, as well as on the time courses of these processes, relative to the HRF. Two important dissociations also occur- 1/ the increase in glucose metabolism is larger than the oxygen consumption rate ($CMRO_2$) with increased neural activity (196); 2/ the level of oxygen presentation far outstrips the rate of aerobic metabolism. The combined effect of this is a net drop in deoxygenated haemoglobin, an increase in oxygenated haemoglobin, and therefore an increase in the BOLD signal strength.

2.2.2.3 A conceptual framework for interactions involved in the BOLD response

Buxton et al. (195) suggested that the BOLD response is driven by a change in CBF, with strong modulation by the level of coupling with $CMRO_2$. They propose the following relationship:

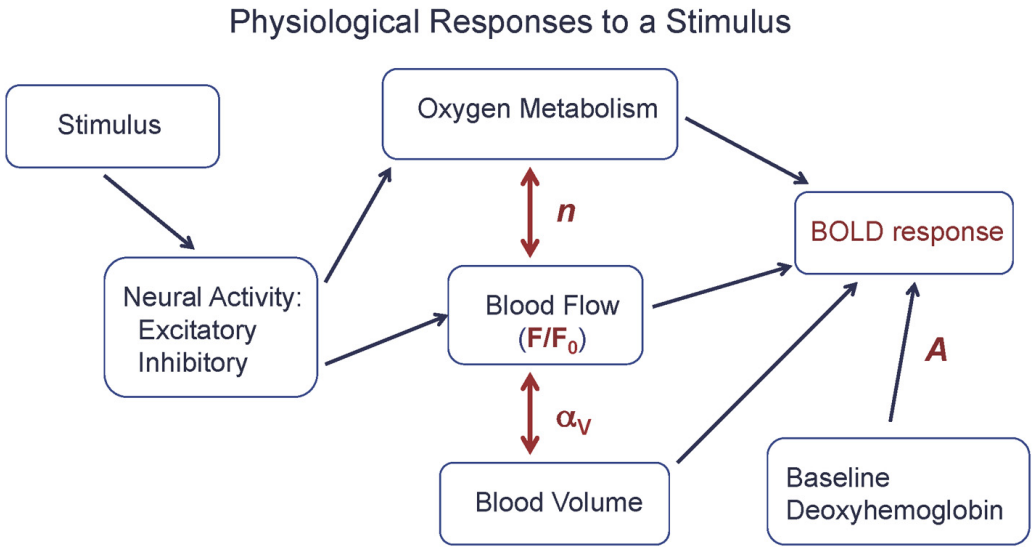

Fig 2-2. The physiological basis of the BOLD response. (See below)

Conceptualisation of the possible relationships underlying the BOLD response, between cerebral blood flow, $CMRO_2$, venous blood volume changes and the baseline haemoglobin state (from Buxton (195).

A = scaling factor, proportional to the total amount of deoxyhaemoglobin in the baseline state

Fo = baseline CBF

F = activated CBF

$\alpha_v F$ = effect of change of venous blood volume with activation

NB: There is nonlinear dependence on F; even with large flow rates, the BOLD response cannot increase beyond a certain point (i.e. there is a fixed amount of deoxygenated haemoglobin).

2.2.3 Where and how is this energy used: glia versus neurone?

The principal use of the energy in neural activity is the restoration of sodium and calcium gradients (183,197). Most of the energy cost is thought to be related to excitatory neural activity, and recovery from synaptic activity, rather than action potential production. Attwell and Iadecola (197) estimated that approximately 74% of the energy budget of the brain is devoted to post-synaptic potentials. Inhibitory neural activity, which primarily relates to the opening of chloride ion channels in relation to gamma-aminobutyric acid (GABA) release, is not thought to incur a significant energy cost.

At the cellular level, neurones and glial cells such as astrocytes engage in activity in response to a functional stimulus. Based on observations, that astrocytes are involved in recycling glutamate released into the synaptic cleft, and $CMRO_2$ correlates with neural activity, a cellular compartmentalisation of activity was proposed. Lower energy producing anaerobic glycolysis was suggested to occur in glia, while higher energy producing oxidative phosphorylation was thought to occur in neurones. This is likely to be a simplification- aerobic metabolism has been

demonstrated with radiolabelled carbon magnetic resonance spectroscopy, occurring in glia; neurons undertake aerobic metabolism during periods of high demand, as is likely to occur in the decoupling of $CMRO_2$/CBF (198).

2.2.4 Why does blood flow increase?

I have established that neural activity is linked or 'coupled' to vascular changes, which produces a characteristic increase in blood flow in relation to increases in neural activity. However, it is not clear what requirement of neural metabolism increased blood flow serves to address. An unanswered question is whether it is a glycaemic or an oxygen demand that is being resolved. Animal models of hyperoxia and hyperglycaemia have demonstrated stimulus-induced functional hyperemia in the relevant parts of the cortex (199,200). Conversely hypoglycaemia in humans did not affect the regional blood flow in response to neural activation (201). It has further been suggested that waste removal and heat regulation may contribute to the witnessed changes in blood flow (202).

The CBF increase provides oxygen levels significantly in excess of that required by active neurones. A number of theories have been put forward to explain this. Gjedde suggested the hyper-oxygenation state helps to maintain a high oxygen tension gradient from blood capillaries to brain tissue, and that this is necessary for active cells (203). It has been proposed that the increase in CBF is not driven directly by energy metabolism, but rather the neural activity itself (197) – this mechanism increases CBF in anticipation of a greater need of oxygen (189), providing a protective energy redundancy.

The delay in the haemodynamic peak relative to the neural activity implies that the changes in blood flow (and volume) are not directly linked to the initial metabolic needs of the firing neurones. One possibility is that the peak may be a consequence of the surge in blood flow that is produced by the initial metabolic demand of the firing neurones. This would mean that the haemodynamic peak is less important than the rising gradient of the peak – the latter relating more directly to neural firing, and the former representing a lag effect. On the other hand, as has been proposed by Heeger and Ress (190), sufficient stores of nutrients may already be available in the neurones or associated support cells to provide for the initial neuronal metabolic demand – this would be essential if the speed of the vascular response was limited.

A later blood flow would then serve to replenish depleted stores, and provide energy (glucose and oxygen) for on-going firing.

One model has attempted to explain both the early requirement for neuronal energy and the delayed haemodynamic response peak utilising a cellular mediator. The 'lactate shuttle' model proposes that the neighbouring astrocytes undergo anaerobic glycolysis in response to elevated glutamate levels, prioritising local oxygen reserves for neuronal activity. The lactate that is produced by anaerobic metabolism is utilised as an energy source by neurones.

2.2.5 What are the cellular determinants of blood flow increase?

Stimulus evoked positive BOLD may be driven by neurovascular coupling at a cellular level, and the properties of the brain's vasculature. Astrocytes are an obvious candidate cell for the interaction between neuronal activity and vascular changes. These glial cells ensheath arterioles, capillaries and ascending venules throughout the cortex (204). The mechanism for coupling metabolism with vascular changes may be linked with the presence of glutamate, a by-product of neuronal metabolism. Metabotropic glutamate receptors (mGluR5s) on astrocytes identify the increase of glutamate, leading to a rise in intracellular calcium, and calcium-activated enzymatic processes. This, in turn, causes the production of local vasodilators (e.g. arachidonic acid derivatives such as prostaglandins, PGI2) which cause vasodilatation through direct action on perivascular smooth muscle. However, a number of inconsistences in the glutamate-mediated astrocyte activation pathway remain unaddressed. These include the potential absence of mGlusR5s receptors in the adult brain (205), the lack of astrocytes on pial arteries on the cortical surface, and the observation of functional hyperemia dissociated from intracellular calcium changes in mice models (206,207).

The capillary bed forms the primary vascular-tissue interface for metabolic activity in various parts of the body, including the brain. Therefore, despite the identified discrepancies between capillary blood flow, and those linked to BOLD changes, it is

relevant to consider its potential role in neurovascular coupling (191). Most neurons are closer to capillaries (8.4μm in hippocampus) than to arterioles (70 μm), leading Hall et al. to propose neurons may adjust their energy supply by initially signalling to pericytes. Capillary pericytes are contractile cells, which wrap around the endothelium, and regulate blood flow, phagocytosis of cellular debris, and the permeability of the blood-brain barrier (208). Capillaries dilate more rapidly than arterioles, facilitating their dilatation - pericytes actively relax to generate capillary dilatation. Glutamate once again acts as the signal, causing the generation of PGE2, which dilates the capillaries by activating an outward K^+ current in pericytes. This hyperpolarising current may serve as the vasodilatory signal, passing to the arterioles via gap junctions (209,210). Hall et al. demonstrated in vivo findings that pericytes are the first vascular elements to dilate during neuronal activity (1s before penetrating arterioles), being responsible for 84% of the increase in blood flow generated by neuronal activity. This provides compelling evidence for 'pericyte-mediated capillary dilatation' being responsible for the BOLD signal.

Rapid propagated vasodilatation may provide a further local mechanism for neurovascular coupling (211,212). It has been shown to occur in the cortex during functional hyperemia (213–215), and is mediated via a rapid endothelial hyperpolarisation. This propagates electrically within the endothelium (distances > 10mm) – it causes self-dilatation via myoendothelial coupling to encircling smooth muscle cells, via gap junctions in the myoendothelium, and potentially a nitrogen oxide/ cyclo-oxygenase independent pathway utilising a putative endothelium-derived hyperpolarising factor.

Brain regions and networks have been proposed as being implicated in controlling and constraining local neurovascular coupling. Cholinergic afferents from the basal forebrain have been shown to modulate regional blood flow (207,216–218). Paracrine effects in association with local and distributed neuronal networks may provide for local coupling (219,220). This may involve fine-tuning of local neuronal vasculature, with interneurons acting on pericytes (197,221). Vasoactive substances released by cortical interneurons have been demonstrated, including Vasoactive intestinal Peptide, Nitrous Oxide, Neuropeptide Y, and somatostatin (222). At the macroscopic end of the spectrum, norepinephrine release by the locus coerulus

causing wider cerebral vasoconstriction, and has been proposed to modulate functionally induced hyperemia (223,224).

2.2.6 An integrated model

In a recent review article (225), Hillman suggested that the most parsimonious explanation of the BOLD response was explained by incorporating propagated vasodilatation, mediated at the endothelial level, incorporating EDHF and NO mechanisms (225). She further hypothesises that the fast and slow components of propagated vasodilatation provide a physical basis for the non-linearities of the BOLD response (e.g. influences of stimulus frequencies, amplitude, and anaesthesia on the HRF). This dovetails well with the documented role of capillary endothelial pericytes (191) (see above), and paves the way for a cohesive view of the cellular basis of functional hyperemia, underpinning the BOLD effect.

2.2.7 How does electrical activity match up to BOLD imaging

Neural activity can separated into single neuron activity, multiple neurone spiking activity and synaptic activity. In electrophysiological terms, high frequency spiking activity is measured using multi-unit spiking activity (MUA). MUA is thought to represent the average spiking of small neuronal populations close to the vicinity of the placed microelectrode. It is obtained by band-pass filtering the recorded signal in a frequency range of 400 to a few thousand Hz. On the other hand low frequency range of neural activity, i.e. local field potentials (LFPs), reflect synaptic activity, being the synchronised synaptic inputs of a given neural population. An LFP is an electrophysiological signal generated by the summed electric current flowing from multiple nearby neurons within a small volume of nervous tissue. Voltage is produced across the local extracellular space by action potentials and graded potentials in neurons in the area, and varies as a result of synaptic activity. "Potential" refers to electrical potential, or voltage, and particularly to voltage recorded with a microelectrode embedded within neuronal tissue Utilising combined

MR imaging and electrophysiology experiments, Logothetis confirmed that LFPs are the most reliable indicator of BOLD activity (226). From a neuronal perspective, the BOLD signal primarily reflects incoming specific or association inputs to a brain region, as well as processing by both excitatory and inhibitory interneurons (194,226–228).

A further consideration is that the BOLD signal may become uncoupled from electrophysiological responses in superficial cortical layers (i.e. demonstrated using comparative measures of LFP and MUA, versus depth electrode measurements in rat somatosensory cortex, (229)).

2.2.8 What is the spatial resolution of the BOLD signal?

This is a complex issue, with the BOLD response representing various neural and vascular determinants, which occur on different timescales. It bears reiterating that it is fundamentally a haemodynamic measure, and surrogate of the underlying neural activity. At a macroscopic level, the BOLD signal may be limited by microvascular density, which is less than the neural mass. The influence of large draining veins is another confounding factor, which is in part addressed by spin-echo techniques. Currently, it has been suggested that the resolution of fMRI BOLD is at the level of one cortical column, which contains 10^5 neurones (230). Lengthening acquisition times may result in increased spatial specificity, with a potential reduction in signal-to-noise ratio. Increased field strength, i.e. 7T scanning, has increased capillary specificity as compared to lower field strengths, and may therefore be weighted towards the metabolic demands of active neural tissue, revealing microvascular change in grey matter (231).

2.2.9 Volitional control of the BOLD signal: training a vascular signal?

This thesis is concerned with the voluntary self-modulation of the BOLD signal, which provides an interesting context in which to consider the underpinnings of the

neurovascular signal. My work, and other neurofeedback studies confirm that volitional control of BOLD signals is possible and can be achieved rapidly i.e. within one day, over 30-45 minutes (232). This is in contrast to EEG neurofeedback, which typically requires multiple sessions that corresponds to e.g. several hours of training.

Assuming that global effects such as arousal are accounted for (see rt-fMRI section) the question arises – what exactly are participants learning to control when they self-modulate their regional BOLD signals? The BOLD signal is a fundamentally metabolic signal, which is likely to be modulated by capillary vascular tone (i.e. pericyte-mediated). Given previous studies demonstrating biofeedback mediated control of cardiac vasomotor tone (51), and the proposed involvement of procedural/motor learning processes, it is possible that neurofeedback training is facilitated by control of intracranial vasomotor tone. This might set-up a positive feedback loop between locally increased neural activity (e.g. in response to activation of specific representations) and increased oxygenated blood flow. The latter could be both activity-dependent, and driven by a trained vasomotor response.

A corroborative finding is that changes in the structural plasticity of the brain in response to interventions which purport to stimulate neural growth and reorganisation such as learning, in addition to the increases in neural mass by axonal enlargement and sprouting, are also accompanied by substantial increases in non-neural tissue such as glial cells and vasculature. Indeed, it has been previously argued that vascular changes could occur even in the absence of neurogenesis (233), implicating an important role for dynamic vascular changes in response to learning. The issue of plasticity is central to the motivation and findings of Experiment 4, and I return to it in the discussion of fingdings in which structural changes were produced by neurofeedback training.

2.3 FMRI data analysis

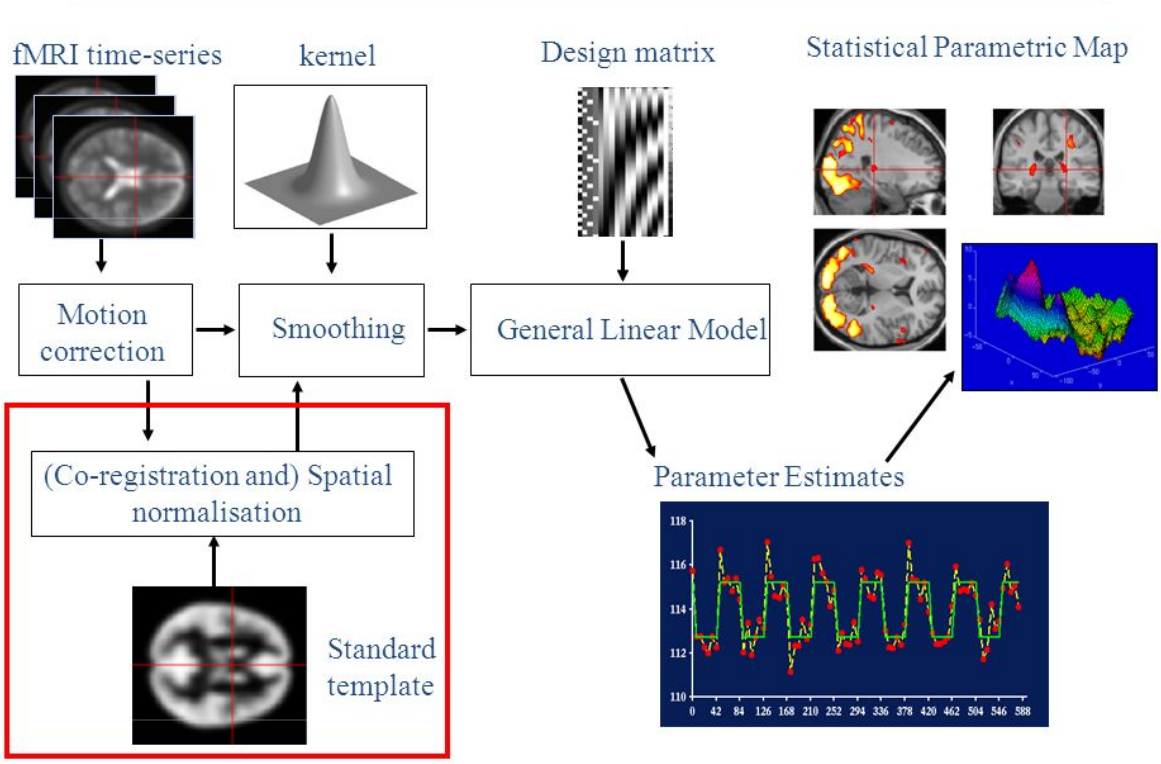

Fig.2-3: Schematic of a typical fMRI preprocessing pipeline. The component of this are discussed in sections on 'FMRI data analysis' and the 'General linear model'.

2.3.1 Preprocessing

This section covers the standard steps, known as preprocessing, that are utilised to prepare fMRI data for further analysis. The major steps are 1/ slice timing correction, 2/ spatial realignment, 3/ co-registration of structural to functional images 4/ spatial normalisation of the functional data to canonical template as necessary, and 5/ spatial smoothing (234).

2.3.2 Slice timing correction

Neuroimaging data, which is often based on the entire three-dimensional volume of the brain being imaged, is acquired in a series of axial slices, with each slice being

made up of an array of voxels. In subsequent data analysis, the assumption is made that the whole brain volume is acquired simultaneously. In reality there are a number of options that can be implemented. The time taken to acquire a brain volume, i.e. TR, is on the order of seconds (typically~2s). In this thesis, slice acquisitions were acquired in an ascending order, starting with caudal slices. However, an additional consideration is if interleaved slice acquisition is used, where odd-numbered slices are collected first, in an ascending order; even-numbered slices are then collected, in a descending order. This means that the BOLD responses in contiguous slices are separated in time, sometimes several seconds apart (235), with implications for voxels in neighbouring slices. In all such cases, slice timing correction enables each voxel's time course to be considered simultaneously. The approach takes two forms. Firstly during preprocessing, temporal interpolation is used. The most common forms are linear, spline (i.e. non-linear) and the use of the cardinal-sine function i.e. sinc interpolation. During task-related analysis, the HRF can be altered to account for moderate differences in timing. Popular methods include using a weighted sum of two terms, where one term is the standard HRF (e.g. sum of two gamma functions), and the other is temporal derivative of that same HRF (236).

2.3.3 Realignment (Motion Correction)

Minor head movements during scanning can significantly corrupt imaging data. For example a movement of 5 mm may increase 'activation' values in a voxel by a factor of 5 (237). Indeed if head movements are too great, participants may need to be removed from the study altogether. The aim of realignment is to correct any small motion changes across a series of functional images by aligning them to a reference image, such that each voxel will have the same co-ordinate throughout a time series, and all images will have the same orientation. The first image is chosen as the reference image, having discarded the first 5 images to allow for field equilibration. Every subsequent image will be registered and re-sampled to be in the same orientation as the first image. Six parameters are typically used for a rigid-body transformation (i.e. maintaining the size of the brain as a constant) of the data into the first functional image. The three rotation parameters are along the x, y and z-

axis, and the three translation parameters are left-right ‘roll’, up-down ‘pitch’, and forward-backward ‘yaw’. These values are calculated iteratively using an optimisation algorithm to minimise the sum of the square of the differences between the reference image and each subsequent image (238). It is occasionally necessary to manually correct for differences between images, which may still occur, particularly in the context of multiple session/ day scanning. This situation occurs with rt-fMRI neurofeedback studies, which may have as many as 6 scanning days per participant.

Non-linear movement artefacts may persist, as a consequence of the position of each voxel in the B_0. The image intensity depends on the spin-excitation history (i.e. previous positions of the brain volume in the field) and movements between slice acquisitions. Correction of persistent movement artefacts may be achieved during the model estimation, by regressing the estimated movement parameters from the specified data matrix i.e. as specified in the general linear model (239). This was the approach used for univariate analyses in this thesis.

2.3.4 Unwarping and the use of field maps

In order to reduce the distortions in image quality produced by susceptibility to field inhomogeneities, a further sequence of imaging is performed which can be used to weight and exclude field deviations. These inhomogeneities are most pronounced at brain tissue/air interfaces e.g. air sinuses.

The field map acquires two T2* gradient echo images, with each image having different TE times to produce different weightings. A field map is then calculated based on the phase difference of the signal at each voxel, relative to the 3D field variation in that voxel. This can be used to work out how much to ‘un-distort’ each voxel i.e. how much to the spatially shift each voxel, resulting in the production of a voxel displacement map. An inversion of this is applied to the EPI images to produce a ‘corrected’ image (240).

2.3.5 Spatial normalisation

Normalisation is used to warp whole brain imaging data from several participants into a standardised anatomical space, utilising established templates such as the Montreal Neurological Institute template (241), and the Talairach space (242). This allows for group level statistical tests to be applied at specific anatomical co-ordinates. Although this is not essential for group level analysis, it is useful for alleviating inter-participant anatomical variance, particularly in the context of complex modelling, and more recently, real-time functional localisation (243).

Normalisation consists of two steps. The determination of optimum 12-parameters, consisting of 3 translations, 3 rotations, 3 zooms and 3 shears, to affine transform the mean image of the individual participant created during realignment, to the anatomical standard template (see Fig 2-4 A). This is then followed by a non-linear estimation of deformations, required to address differences in head/ brain shape not accounted for by the affine transform parameters. In the latter step, the application of a choice of discrete cosine basis functions minimises the residual square of the differences between the images and the template (244).

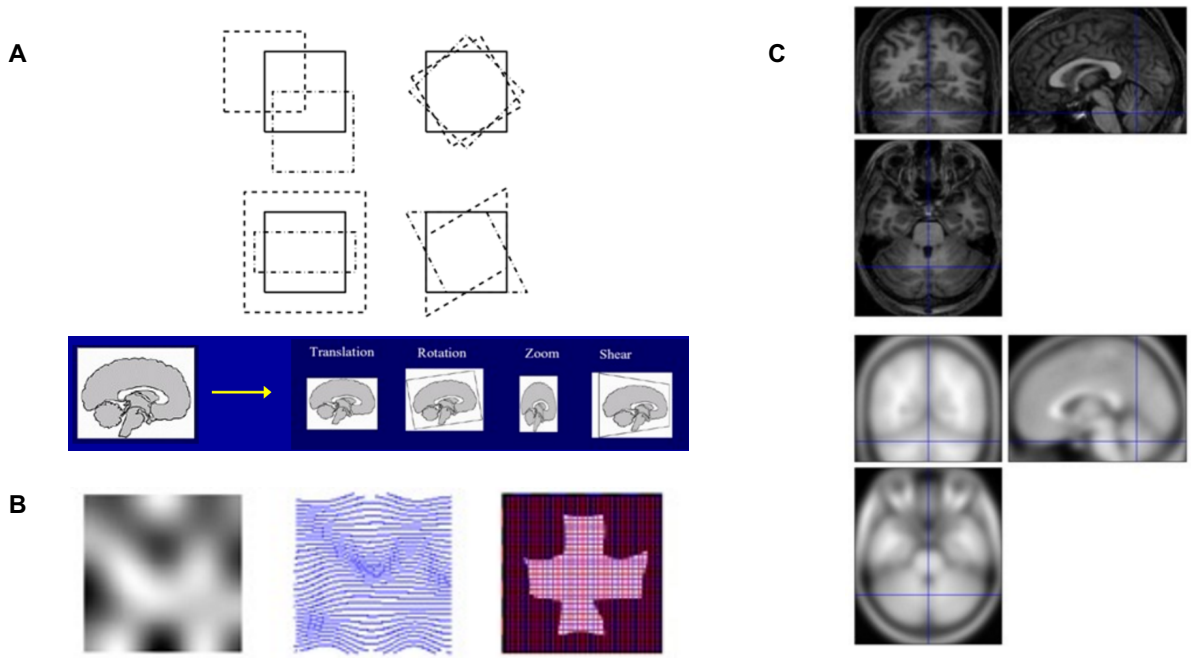

Fig.2-4: Illustration of the mechanics of spatial normalisation. (See below)

A, B Illustration of the 12 affine transformations (3 translations, 3 rotations, 3 zooms, 3 shears), and non-linear transformation (consisting of linear combination of smooth basis functions, obtained from a discrete 3D cosine transform) utilised in spatial normalisation C. Anatomical brain (top) and a functional image (bottom) shown in standard space after spatial normalisation.

The application of the 12 parameter affine transform and warping is constrained within an empirical Bayesian framework (238), using prior knowledge of head shape and size. In addition to the reduction in the square of the differences between the images and the template, the distance between the parameters and the prior expectation is also reduced, which prevents deviation of the transform from its expected value. This is necessary to avoid the introduction of unnecessary warping (i.e. 'over-fitting'), by enabling the 'fit' that is most likely given prior information, as compared to a standard rigid-body transformation where the constraints are explicit.

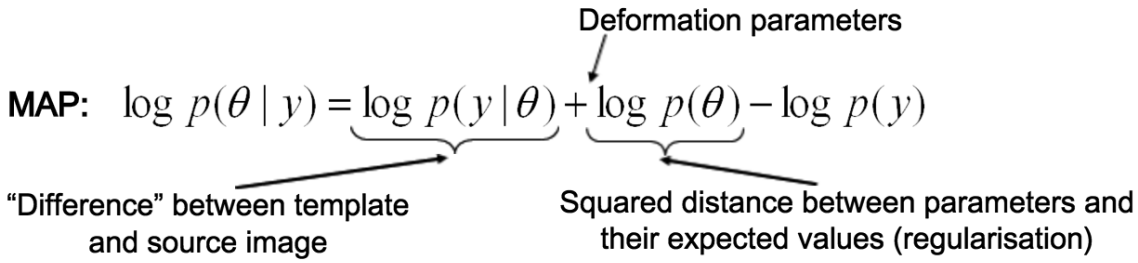

The equation above summaries the Bayesian regularisation step, which makes use of a maximum a posteriori scheme (MAP), and ensures that non-linear normalisation does not introduce any unnecessary warps. A MAP inference is the problem of finding an assignment to the variables of a probabilistic model that maximises their joint posterior probability (245). For spatial normalisation, it is achieved by minimising the sum of the squared differences between the template and the source image, and minimising a function of the deformation field by minimising the squared difference between the deformation parameters and the deformation priors. It is worth noting that normalisation corrects for gross differences, with spatial smoothing being used to address residual variability. The mathematical basis of this process is beyond the scope of this thesis with a full instruction on spatial normalisation found here (244).

2.3.6 Co-registration

High-resolution T1 structural images (1mm isotropic voxels) are used to align a particular participant's functional EPI images (2-3mm isotropic voxels), using a rigid-body transformation to the mean functional image. This is more involved than that used for realignment (described above), due to the differences in image intensities of the structural T1 and the functional T2 sequences respectively. A two-dimensional histogram is created for each image by grouping voxels according to their intensity values. An optimisation algorithm is then applied which relies on mutual information theory - it attempts to match the two histograms such that the amount of uncertainty between any two voxels in the two images is minimised, thereby establishing the appropriate transform parameters.

That said, EPI scans may also be registered to a participant's own 'averaged over time' EPI image – this provides fewer obstacles than those required by intensity correspondence. Although this approach is less spatially specific, it may be more sensitive to detecting activity differences.

2.3.7 Spatial Smoothing

Spatial smoothing is performed prior to data analysis to improve inter-participant registration, and to reduce noise in the BOLD signal by blurring residual anatomical differences. This also serves to ensure assumptions of random field theory are upheld in relation to correction for multiple comparisons i.e. by reducing the number of resolution elements regarded as being independent. In regions that are spatially larger than the imaged spatial resolution, smoothing may reduce random variance in individual voxels, helping to increase the signal-to-noise ratio within the region, by providing a weighted average of the local signal (237).

Spatial smoothing involves convolving the functional data with a Gaussian kernel, described by the full width of the kernel at half its maximum height (FWHM). Typical

values for kernel widths range between 6-12mm FWHM – all of the experiments in this thesis employed a 6mm smoothing kernel.

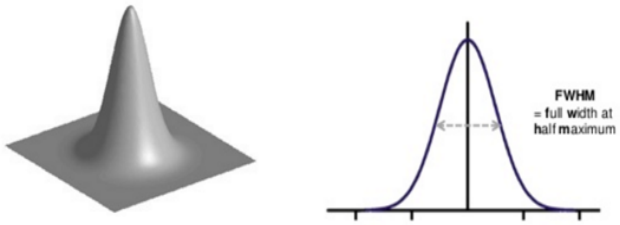

Fig.2-5: Gaussian smoothing kernel. Image of a 3D Gaussian smoothing kernel, which is convolved with the imaging data at each voxel (left hand image). The Gaussian kernel is a normal distribution, with the shape of the kernel being described by the width of the distribution at half its maximum value (right hand image).

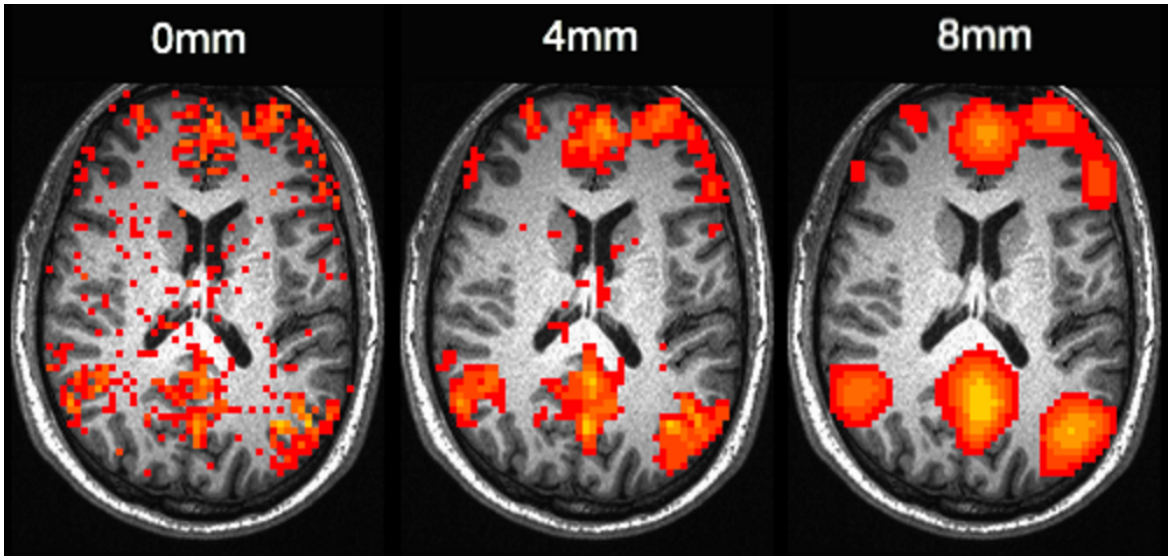

Fig.2-6: Functional statistical maps, with and without smoothing. The application of a smoothing kernel results in the measured BOLD signal at each voxel being replaced by the weighted average of its neighbours. Depending on the size of the kernel used, the BOLD signal is variably smoothed as shown above, with clear differences in the statistical data for smoothing at 4 versus 8 mm.

2.3.8 High Pass filters and low frequency drifts

Voxel timecourses show low frequency fluctuations or 'drift', caused by physiological noise, as well as scanner related fluctuations, such as those caused by minor changes in temperature. These drifts demonstrate characteristic low frequency 'rising' and 'falling' over the course of a task block. This was observed in the experimental work presented in Chapters 4 and 5. Signal drifts substantially reduce the power of statistical analysis. They can be removed using 'frequency filters', with the caveat that condition-related signal change may also be inadvertently removed. The application of a high pass filter is applied separately at each individual voxel – neighbouring voxels may have different rates of drift (235).

In SPM there are default settings for the implementation of the high pass filter, which is set at a cut-off point of 128 Hz. I specifically altered this to 256 Hz in the 'attention decoding' experiments presented in Chapters 4 and 5, to ensure task related signal was not removed.

2.3.9 Grand mean scaling

Grand mean scaling eliminates global differences across sessions or participants, once again due to physiological and scanner-related fluctuations. This step, in a similar manner to the aforementioned filtering, adds a further correction for slower fluctuations in field strength that may occur over hours or days, or in the cerebral blood oxygenation differences between participants. Grand mean scaling is performed routinely in fMRI preprocessing, and normalises the overall mean BOLD signal to the same numerical value, correcting for session-to-session differences. There are other methods which constrain data further, such as proportional scaling - the term for these processes is *global normalisation* (237).

2.4 General linear model

2.4.1 Overview

The general linear model (GLM) is the most commonly used multiple regression model in univariate fMRI data analysis. The fundamental premise is that a prediction informs the design of the model which best explains the experimentally modulated neural activity, taking into account timing and duration of the events; the predictions are weighted in order to minimise the influence of error in the subsequent measurements (239).

The data in this case is contained in a volumetric space, contiguously mapped by tens of thousands of isotropic voxels. Importantly, the model testing is performed on an individual voxel basis – this forms the basis of the univariate approach. The spatial structure of fMRI data is in fact not used in the model. Instead the voxels are treated independently of one another, and are arranged along a single dimension per time point for ease of calculation. An expression to describe this is formulated by the following:

$$\mathbf{Y = BX+e}$$

The *data matrix* is expressed by **Y**- made up of **V** voxels, by **n** time points. The *design matrix,* **X**, identifies the linear model to be evaluated, and is made of the regressors **M**, each **n** time points in length. The weights of each regressor is combined into a parameter matrix **B**. Finally an error matrix, **e**, is included.

Regressors associated with specific experimental hypothesis are termed ‘experimental regressors’. These are usually covariates, which have a continuous range of values, or may be ‘indicators’ which have integral values, and mark a qualitative difference. The latter can be important in session comparisons. A typical model used for neural activation is the ‘boxcar’ function in block-design fMRI, where the response is either zero or one (i.e. present or absent). This is then convolved

with the HRF. The HRF is modelled in SPM, using a multivariate Taylor expansion of a combination of gamma functions (246).

$$h(t) = A\left(\frac{t^{\alpha_1-1}\beta_1^{\alpha_1}e^{-\beta_1 t}}{\Gamma_{(\alpha_1)}} - \frac{t^{\alpha_2-1}\beta_2^{\alpha_1}e^{-\beta_2 t}}{\Gamma_{(\alpha_2)}}\right)$$

t references time, $\alpha_1 = 6$, $\alpha_2 = 16$, $\beta_1 = \beta_2 = 1$ and $c = 1/6$. Γ represents the gamma function, which acts as a normalizing parameter. The unknown parameter in the model is the amplitude *A*. This function is convolved with the stimulus function to obtain a task related regressor to enter into the design matrix. In order to accommodate small jitter that may exist in the onset and duration of responses, a temporal derivative or a dispersion derivative, respectively, may be included in the model (247,248).

An important consideration in the construction of regressors is to avoid their 'co-linearity', in which they are not able to be statistically disambiguated. This can be achieved by ensuring that each regressor has minimal overlap with another, either by introducing large temporal gaps, as for slow event-related designs ,or for rapid event-related designs, by ordering conditions such that they can be deconvolved.

Examples of the importance of producing optimised regressors in this thesis, are explored in the experiments presented in Chapters 4, and 5. Multiple streams of stimuli were presented concurrently in each of four quadrants on a visual display, with the participant being required to direct their attention to one stimulus steam per session block. In order to efficiently model the neural activity related to the deployment of spatial attention in these two experiments, the regressors for each quadrant included quadrant- specific information. This was partly achieved by the spatial location of the visual information being modelled, as well as the category of visual stimuli used i.e. house, faces etc. In block design fMRI, individual events within a task block are not taken into account. For the purpose of the experiments in Chapter 4, and 5, within a block each stimulus presentation would represent an event, with each one lasting for 500 ms – too rapid to be identified individually in a

standard block design. Nonetheless, the stimulus presentations for each quadrant were interspersed with a set number of blank stimuli randomly interspersed with the stimulus presentations, such that attention to each quadrant-specific stimulus stream would produce quadrant-specific neural activity with distinguishable haemodynamic responses. The optimal presentation of these blank stimuli was achieved through the use of m-sequences (maximum length shift-register sequences). M-sequences specify an optimal order of events that allows for zero autocorrelation and therefore maximal statistical efficiency for disambiguating different stimulus events presented close together in time. This is achieved in two ways: m-sequences are nearly orthogonal to cyclically time-shifted version of themselves, and they maximize the number of presentations for all event types (249).

Nuisance regressors, which model known non-experimental sources of variability, are also included in the design matrix. These may be general e.g. movement related to respiration or specific to an experiment e.g. change in day, movement related to a button press. Commonly, six movement parameters are included, and are used to remove variance in the data associated with the three directions of translation and three axes of rotation (237). As with many aspects of experimental design and analyses, there are considerations with respect to the inclusion of nuisance regressors, and they may also be excluded. On one hand, nuisance regressors reduce the amount of residual variation, and improve the validity of the GLM. On the other hand, experimental events may correlate with movement e.g. movement related to a button press may be correlated with neural activity of interest, occurring at the same time. As such the motion regressor would remove from the BOLD signal, activity of interest as well as activity related to movement. Finally, as introduced earlier, temporal filters are used to remove physical and physiological artefacts prior to fitting of the model.

2.4.2 Parameter estimates, T and Z values

The solution to a GLM are parameter estimates (also known as beta values) for regressors of interest. Converting the parameter estimate into a useful statistic

requires taking into account the level of uncertainty in its estimation – this is an important step (250). A commonly used test stastic in fMRI is the t-value, which is calculated by dividing the parameter estimate by the standard error i.e. the residual error. Under the null hypothesis, the error follows a statistical distribution called the F distribution, and it can be evaluated as a function of the available degrees of freedom (which are effected by the number of time points and the number of regressors). Therefore, if the value of the parameter estimate is low as compared to its uncertainty/ error, the fit is unlikely to be significant. The t-value provides a measure of whether the activation recorded is 'real'. These transactions are performed at a voxel level.

The following is utilised to calculate the t- value:

t = explained variance ÷ unexplained variance

t is taken from a t-distribution with *degrees of freedom* equal to the number of images *minus* the *β parameter estimates*.

$$\mathbf{t} = \frac{\mathbf{c}' \boldsymbol{\beta}}{\mathbf{SEM}(\mathbf{c}' \boldsymbol{\beta})}$$

c = vector of the contrast of β parameter estimates
SEM = standard error of the mean

To convert this into a probability p, a Z statistic[2] requires a further transformation – p, t and Z values all provide a measure of how well the data fits the model at a specific voxel co-ordinate.

2.4.3 Contrasts

In order to compare parameter estimates (e.g. towards the relevance of a particular hypothesis), a 'contrast' of the parameter estimates in question is created. One parameter estimate can be subtracted from another, the standard error for the new value is calculated, and a new t-value image is created.

A simple contrast is prepared by denoting a 1 for one parameter estimate of interest, and a -1 for the *other* parameter estimate. This answers the question of where the stimulus response is greatest for the first parameter estimate *relative* to the other. The contrasted parameter estimate, expressed as [1 -1] provides an assessment of the linear interaction of the two individual parameter estimates. More complex contrasts can be created, to model for example, non-linear interactions by the addition of an explicit interaction variable, or to model graduated changes in the data i.e. parametric contrasts. In Chapter 6, parametric contrasts were used initially to interrogate the data to look for the effect of training across multiple training days.

2.4.4 'Thresholding', inference and the problem of multiple comparisons

The resulting probability map produced by calculating voxelwise t-values (or Z values) will need to be thresholded in order to best infer which parts of the data matrix (i.e. representing the whole brain) are significantly associated with the experimental manipulation. This is done by applying a significance threshold e.g. a specific p-value, across the map. The multiple comparisons problem is due to the brain volume being composed of approximately 20,000 voxels. In this situation, using a p-value threshold of < 0.01 would result in a chance activation of 200 voxels (the size of many brain regions of interest), even in the absence of stimulation (189).

The most stringent method of addressing this problem of 'multiple comparisons' is the Bonferroni correction. This divides the significance level at each voxel by the number of voxels – this is considered a conservative approach, and can result in a loss of sensitivity. It was not used in this thesis. SPM solves the multiple

comparisons problem through the application of Gaussian Random Field theory (RFT). The spatial smoothness of the statistical map is considered. The spatial correlation of image intensities results in the number of independent observations (resels) being considered rather than the number of voxels. The number of resels in the image are estimated, and the probability of their exceeding the requisite threshold e.g. $p < 0.05$ is calculated under the null hypothesis (251).

An alternative approach is the False Discovery Rate (FDR). This establishes the proportion of positive results that are false positives; its principal advantages are that it uses a less stringent correction (increasing alpha power) and controls for the proportion of false claims i.e. errors in activation clusters, rather than false tests, that is, errors in statistical tests performed for each voxel. The disadvantage of using an FDR approach is an increase in Type 1 errors i.e. false positives, which is balanced against the increase in experimental power (250).

This leads to the consideration of cluster-level inference, which is an extension of the RFT-based approach. Here the spatial extent of the activations is taken into account, prior to estimating the significance threshold. ‘Clusters’ of voxels are selected on the basis of an arbitrary initial threshold, and assigned a p-value, which may or may not pass the final test of significance. An important point is that the probability can be expressed as a function of the cluster size, allowing for the threshold to be set in relation to a minimum cluster size, as well as the initial cluster -forming threshold. Cluster-size thresholding works, because the cluster size increases more slowly than the probability that a cluster is significantly active. This approach can be more sensitive than voxel-based approaches, although the initial thresholding needs to be carefully applied (250).

2.4.5 Random effects and fixed effects analyses

Specific considerations are required when applying statistical tests to multiple sessions, and/or multiple participant imaging data. Provided the data has been aligned using spatial intensity normalisation to a standardised anatomical template

(or group ‘mean’ template), analyses can be performed at the ‘second level’, considering individual data as a group. There are two main approaches to this, termed ‘fixed’ and ‘random’ effects analyses.

Fixed effects analyses assume that all participants within a group during a scanning session will have similar levels of brain activations (i.e. ‘fixed’); its primary concern is therefore to identify within-session errors. ‘Random effects’ analyses consider between-session errors, and allow inferences to be made about the population from which the participants are drawn.

‘Mixed effects’ analyses combine aspects of both approaches. A further point is that ‘random effects analyses’ of multiple participants will incorporate individual participants data that have treated independent variables as having fixed effects at earlier stages of the analyses – it is often considered under the bracket of mixed effect analyses. For the studies in this thesis, mixed effects group level analyses were performed on the fMRI data. This approach accounts for inter-scan error variance and inter-participant variance. This therefore allows for inference to the population from which the participants were drawn (252,255).

A final consideration is the number of participants required for statistical inferences to produce valid and significant results in the context of imaging studies. There are a number of factors to be considered, including scanner characteristics, brain regions, functional response to the experimental stimulus, and inter-participant variability. The number of participants required is currently considered to be between 16-32, with higher numbers being desirable (252–254). However this issue is controversial, and a power calculation can be conducted to determine appropriate sample size.

2.4.6 Assumptions of the general linear model

It is worth stating the explicit assumptions of the GLM (235):-

- BOLD signal arises from a linear, time invariant system.

- Use of the same design matrix throughout the brain and also regions of interest (ROIs).
- Noise varies with a normal distribution that has similar properties at all time-points i.e. does not vary with time, is independent of the experimental task.
- Data is homoscedastic
- All voxels represent independent statistical tests
- All time points are independent of the others

These assumptions, as they relate to the application of the GLM to fMRI analysis, have been validated in the landmark paper by Friston et al. (239).

2.5 Multivariate analyses

2.5.1 Overview

The BOLD signal of an imaged voxel approximately $3mm^3$ in size reflects the averaged activity of hundreds of thousands of neurons, which may be heterogeneous with regards to the task under study. This 'blurring' of neural information is increased further by the application of traditional univariate analysis techniques that spatially average across only those voxels that contain responses that are statistically significant with respect to the experimental task. While this has the advantage of reducing 'noise' it also reduces 'signal' by ignoring the fine-grained spatial patterns across voxels and by excluding voxels with weaker (i.e. non-significant) but potentially contributory information. This may be important when examining neuronal populations that intrinsically encode topographical space in a contiguous manner.

The use of multivariate analysis on a voxel-wise basis seeks to avoid these concerns by using pattern classification techniques to examine the spatial patterns of responses across multiple voxels, including those which have non-statistically significant responses to the task when analysed with a univariate approach. When taken on aggregate, these otherwise non-informative voxels may provide useful information in relation to task-related brain activity.

Multivariate analyses employ machine-learning techniques including classification algorithms to examine distributed patterns of brain activations. Multivariate pattern analysis (MVPA) together with machine learning algorithms have been widely used to examine the representational content in patterns of neural activity (38,256–263).

2.5.2 'Classification' and 'classifiers'

Classification in the context of neuroimaging data analysis, refers to the process of predicting parts of the GLM design matrix, such that a relationship between discrete

parts of the imaging data can be established with specific behaviours or task performance by the participant undergoing the imaging experiment. ‘Multivariate’ classification is performed on groups of voxels, rather than treating voxels independently, therefore utilising the spatial relationships in the imaging data.

Prior to defining what a classifier is, it is useful to review the terms describing the preparation of data for the application of classification analysis. The groups of voxels are gathered together to produce specific predictors, which are thought to be linked to a particular part of the imaging experiment. These predictors are independent variables termed ‘features’, and the dependent variable of the imaging paradigm is termed a ‘class’. Each ‘class’ is then given a ‘label’. A classifier will take various features (i.e. independent variables) as an ‘example’, (the set of independent variables), and predict the class that the example belongs to (the dependent variable i.e. the stimulus used to generate task-based BOLD activations). The classifier will have a set of parameters that are ‘learned’ from the training data effectively building a model defining the relationship between the features and the class ‘label’ in the training set. The trained classifier can then determine whether the voxels being tested, e.g. from a ROI, have information about the class of the example. The classifier is used to the test the feature/class relationship by applying its function to different sets of examples, the ‘test data’. The latter assumes that the training and test examples are being independently drawn from the example distribution.

Prior to training the classifiers, the fMRI data must be transformed into examples. There are a number of ways of doing this, ranging from using the activity signal itself from one or more voxels, to utilising the linear model activity estimate from the GLM to predict each voxel- so called ‘single voxel responses’. These maybe beta or t-values for a given condition.

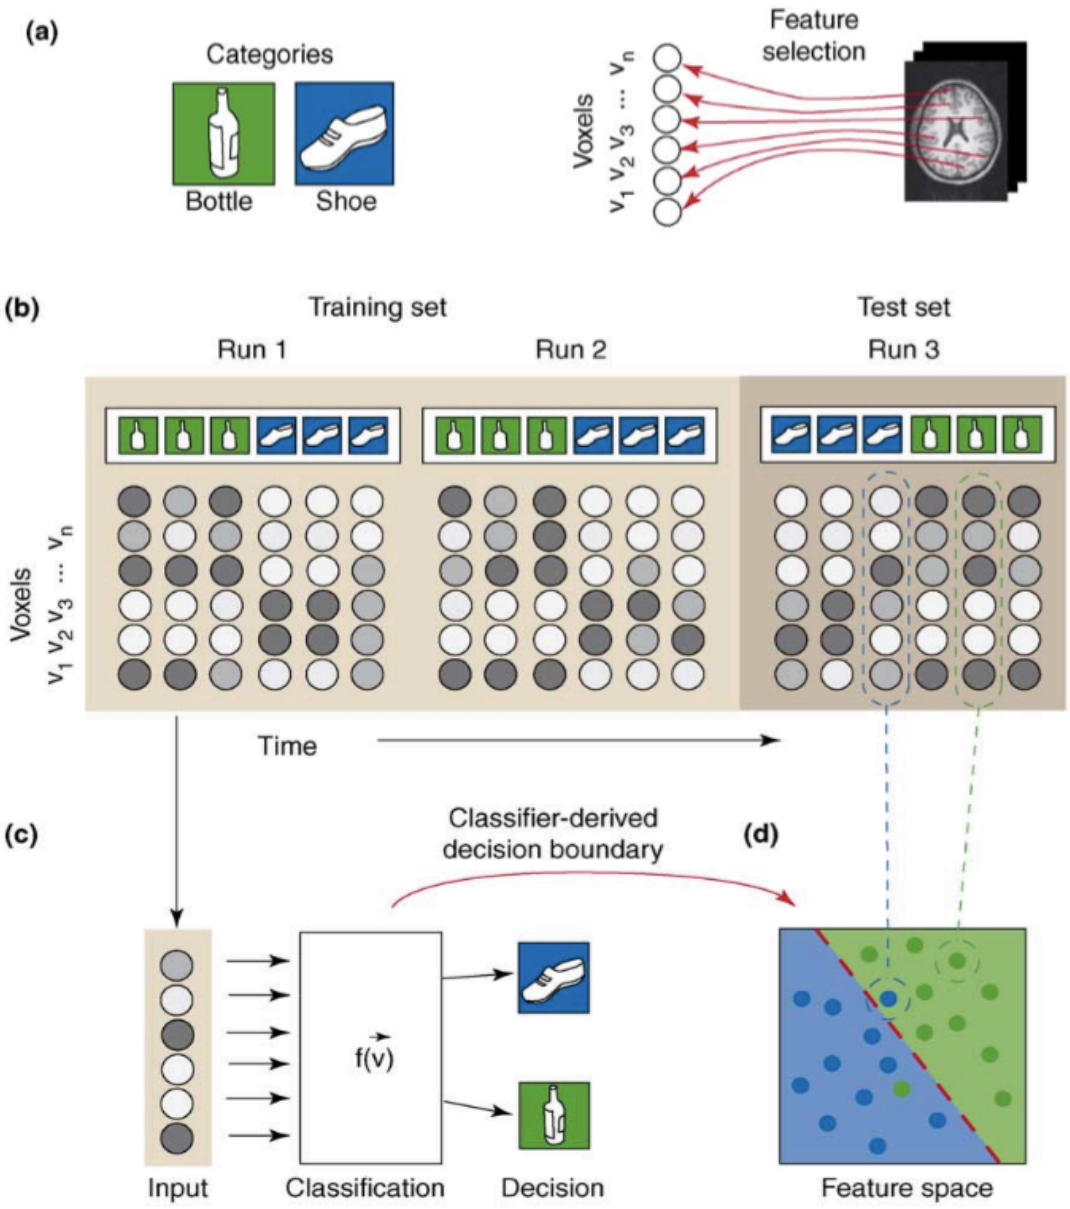

Fig 2-7. FMRI data analysis using multivariate classification. a)+b) A fMRI time series is separated into discrete patterns of brain activity across the selected voxels (i.e. 'features') at a specific time point. Each brain pattern is labelled with one of two experimental conditions and divided into a 'training' and 'testing' set. c) Patterns from the training set are used to train a classifier that maps the brain patterns to each of the two experimental conditions. d) The trained classifier defines a decision boundary (red dashed line) in a high-dimensional space of voxel patterns. Each dot corresponds to a pattern. The background colour of the figure corresponds to the guess the classifier makes for patterns in that region. The trained classifier is then used to predict category membership for patterns in the test set (264).

2.5.3 Feature selection

Feature selection involves determining which voxels will be used in the classification analysis. The principle is to select voxels or groups of voxels that provide the most information with regards to the 'decision boundary'. This is an important aspect of the multivariate technique, as it allows inclusion of voxels that do not meet the conventional univariatecriteria for statistical significance.

The other factor to be considered is avoiding inclusion of voxels with some information, but comparatively higher levels of noise. This prevents the phenomenon of 'over-fitting'- in which essentially a multiple comparisons problem in that it may be possible to find a classification function that accurately classifies the examples in the training set, without it being effective when applied to the test data. Simple linear classifiers also have less tendency to over-fit the data (265,266)

There are a number of established techniques for ensuring the appropriate type (in terms of what they represent) and amount of voxels are selected for multivariate classification analyses. I introduce the key concepts utilised in this thesis, and provide further detail in the experimental chapters for Experiments 1, 2, and 3.

i) Regions of Interest

Rather than looking at all the voxels in the brain, analyses can be limited to those in a specific ROI. Anatomical ROIs or functional ROIs e.g. voxels that demonstrated activity (threshold) during task trials relative to the baseline (as modelled by rest) may be used to select voxels in this way.

ii) Dimensionality Reduction

Dimensionality reduction is another approach to selecting relevant features for a multivariate classification analysis. This involves reducing the original feature space into a new, lower dimensional feature space, yielding a dataset matrix with the same number of rows but a reduced number of columns. A typical approach is to apply principal component analysis (PCA) to whole-brain fMRI data prior to classifier training (267).

iii) Ranking

This is an application of the 'scoring/filtering' method involving ranking each of the features by a given criterion. Each feature is scored in an analogous way to univariate tests of a voxel. The 'best' voxels are subsequently selected.

iv) Cross-validation

To establish a useful estimate of a classifier's accuracy, the optimal situation is to train and test on as many examples as are available (assuming ideally that the classes of data are balanced, and that the training data a each fold contains examples of all classes. In this regard the method of 'leave-one-out' cross-validation is useful, as it produces as many classifiers as there are examples (268).

2.5.4 Types of 'Classifier'

All classifiers have in common the use of a 'training data set' to define a decision boundary in the space of response patterns i.e. the space spanned by the activity levels of the voxels in the ROI (269). Classifiers differ in the shapes they allow for the decision boundary (e.g. hyperplanes in linear classifiers; more complex non-planar boundaries in nonlinear classifiers) and in the way the boundary is placed on the basis of the training data. For example, the Linear Discriminant Analysis (LDA) classifier places the decision hyperplane so as to optimally discriminate two equal-covariance Gaussians; a linear support vector machine (SVM) places a decision hyperplane so as to maximise the margin separating patterns on either side. Each classifier has certain strengths and weaknesses in relation to the shape of the decision boundary e.g. more flexible decision boundaries lead to better separation of data, but introduce the risk of over fitting (266). Further, each classifier must be considered with respect to the brain region being studied, the use of ROIs, single voxel response estimates (beta estimate or t-value) and categorical dichotomy. There is recent empirical evidence to suggest that linear classifiers perform better than non-linear classifiers, with some preference for LDA and pattern correlational

classifiers on account of simplicity, robustness, interpretability and ease of application.

The two commonly used simple linear classifiers, as successfully implemented in this thesis, are the Correlational classifier and Linear Discriminant Analysis (LDA) classifier with PCA (270). They are robust, and reliable, with relatively rigid hyperplanes, for ease of application, comparison and avoidance of over-fitting.

1. Correlational Classifier

For each task condition the mean is calculated across all runs in the training set. To decode using the classifier, a linear correlation is calculated between each sample in the test set and the mean samples from the training set. A test sample is then assigned to the condition that produces the greatest correlation coefficient.

2. Linear Discriminant Analysis

A LDA classifier determines the discriminant dimension response pattern space, on which the ratio of between-class over within-class variance of the data is maximised (269,271). After projection of the data into the linear discriminant dimension, a classification threshold is placed at the midpoint between the two class means. The classifier is Bayes-optimal in that it ignores estimation error, and assumes Gaussian within-class distributions.

In a two-class classification problem, the normal vector of the decision boundary is estimated as a product of the within-class covariance matrix and the mean vectors of each class. Depending on which side of the hyperplane a discriminant value falls then determines which of the two classes it belongs to (i.e. after establishing its distance from the hyperplane in relation to the mean value of all samples, and the vector of voxel values, whether it is a positive or negative integer).

2.5.5 Classifier performance

The performance of a classifier is determined by how accurately it correctly assigns a class label to examples of the imaging data, represented by features, which are selected to best discriminate the data in relation to the experimental conditions. A common metric of classifier performance in fMRI to assess for mutual information between experimental conditions and regional fMRI response patterns is termed ‘decoding’. Above-chance ‘decoding accuracy’ with any classifier, indicates that the functional response patterns contain information about the experimental condition or stimulus category being tested for (269).

The relative performance of different classifiers on fMRI data is incompletely understood. Pereira et al (272) attempted to provide a roadmap of the potential choices for the preparation of data for use in classification, and the relative benefits of specific classifiers. Misaki et al(269) extended this work comparing more classifiers, including all of those used in Chapter 4 of this thesis, such as simple linear classifiers such as the correlation classifier, and LDA, as well as more complex non-linear classifiers such as K –nearest neighbour and radial support vector machines. The study validated the use of ROIs for voxel selection, t-values rather than betas as response estimates, and leave-one- run out cross validation as an optimised approach to ensure better overall classifier performance. Linear classifiers such as LDA and correlational classifier had the best performance, in keeping with previous work (257,273). FMRI neuroimaging data has been suggested to be suited to linear classification, with there being insufficient data for the efficient application of non-linear decision boundaries.

A practical consideration of classifier performance is the purpose to which it will be applied. The use of multivariate classification in the experiments in Chapters 4 and 5 was to explore the possibility of accurately decoding the deployment of spatial attention, for use in a non-invasive, attention-driven BCI for communication(15,274). The current recommendation for operational accuracy of a BCI in order for it to be of utility to patients has been proposed to be 70-80% (78), although the interaction of cognitive task on classification accuracy as well as, individual differences in brain activation responses, and imaging modality remains to be established.

2.6 Retinotopic mapping of the visual cortex

2.6.1 Introduction

In view of the biological and evolutionary importance of the spatial arrangement of physical features in a visual scene (e.g. a tiger jumping down, off a tree on the left) it is not surprising that topographical representations of the world are represented at multiple locations in the brain.

To begin with, information encoded by light entering the eye stimulates neurones on the retina in a spatiotopic fashion. More specifically, the receptive fields of the retinal output neurons (i.e. ganglion cells) form an orderly mosaic of the visual hemi-field in each eye. The resulting space-to-retina mapping is preserved in part in the axons of the optic nerve (275,276), and is fully represented along the visual pathway, through the lateral geniculate body, to the primary visual cortex. The presence of retinotopic maps in primary visual cortex, or primary retinotopic cortex was first documented by turn-of-the-century lesion studies (277–280).

Structural MRI enabled these early observations in humans to be unified with a large body of animal work confirming the presence of retinotopic maps in visual cortex (281). However a systematic approach for mapping retinotopic cortex *a priori* and non-invasively, in the absence of lesions of the visual cortex, was formalised far more recently, using fMRI (282,283). Sereno et al. (41,284) have subsequently demonstrated the presence of retinotopic maps throughout extrastriate visual areas, in higher order visual cortex such as the IPS, lateral occipital cortex and ventral visual areas. I shall discuss the first of the approaches, phase encoded retinotopic mapping, which was used in the experimental work reported in Chapter 3. This technique has evolved further following on from the studies presented in this thesis, into ‘population receptive field mapping’, a technique that has not been used in this thesis, and therefore will not be discussed further. I used retinotopic mapping in order to create retinotopic ROIs. Given the spatiotopic organisation of the retina, and the corresponding cortical retinotopic maps, I explored the degree to which spatial

information can be decoded from a retinotopic region relative to higher order visual areas.

2.6.2 Phase-encoded mapping and cortical flattening

Phase-encoded retinotopic mapping is a reliable approach to systematically defining visually responsive brain areas. The mapping procedure requires participants are required to maintain central eye fixation while viewing dynamic high contrast stimuli. The stimuli travel across the visual field in a phase-controlled cycle, so as to create a travelling wave of excitation across the visual cortex. Two stimuli are used – the first is an expanding ring containing high visual contrast e.g. a flickering checkerboard. This allows measurement of eccentricity i.e. distance from the centre of gaze (which itself will map to the centre of the macula). The second stimulus is a rotating wedge, with its tip at the centre; this enables measurement of the polar angle i.e. orientation with respect to the centre of gaze. Stimulus presentations are repeated multiple times, in both directions (e.g. expansion and contraction for the first stimulus, clockwise/anticlockwise for the second stimulus).

The phases of stimuli presentation will ensure different regions of the visual field are stimulated, with different polar angles and eccentricities occurring during separate phases. Activated voxels will show peaks corresponding to the cycling frequency of the stimulus. The phase lag of these responses will measure the spatial preference of the voxels. The responses are convolved with the haemodynamic lag, averaged across sessions, and specified for cycle directions (i.e. expansion vs. contraction for the first stimulus, clockwise vs. anticlockwise for the second stimulus). A fast Fourier transform is applied to the time course of each voxel in order to determine which voxels are maximally activated at the implemented stimulus frequency, and the size of the phase lag.

2.6.3 Freesurfer visualisation and meridian mapping

For typical three-dimensional rendering of the brain, activations that take place within a sulcus may be obscured, or invisible. To visualise this, software packages exist to inflate, and/or flatten the cortical surface. With regards to retinotopic information this is advantageous, as adjacent points on a cortical surface map to adjacent points in the visual field. Further, the functional activations generated by the contrast for the horizontal and vertical meridians can be displayed on the inflated cortical surface, allowing rapid and intuitive delineation of primary visual cortex. For this, the FreeSurfer software package (http://surfer.nmr.mgh.harvard.edu) developed at Harvard was used in this thesis (285). FreeSurfer automatically performs cortical reconstruction and volumetric segmentation. First, non-brain tissue e.g. skull is removed using a deformation procedure (286), followed by automated Talairach transformation, and segmentation of the subcortical white matter and deep grey matter structures (287,288), intensity normalisation (289) and tessellation of the grey/ white matter boundary. Topology correction (290,291) and surface deformation is performed, following intensity gradients to optimally place the grey/white and grey/CSF borders (285,292,293). A cortical model is generated, which can be manipulated to visualise the sulci e.g. through cortical inflation.

Curvature information is added using intensity and continuity information from the segmentation and deformation procedure (293). Functional data in the form of activations responding to horizontal and vertical meridian stimuli are superimposed onto the inflated cortical surface**.** These are used to delineate the boundaries of the primary visual areas; surface regions are then converted to volumetric ROIs using FreeSurfer. The placement of the meridian lines delineates primary visual regions i.e. V1, V2, V3 on both hemispheres, and both ventral and dorsal surfaces i.e. V1 dorsal, V1 ventral, right and left; similarly for V2 and V3.

2.7 DARTEL analysis and Tensor Based Morphometry

In Chapter 6, I used a combination of two novel techniques to investigate longitudinal structural changes in pre-specified brain regions. In order to conduct a structural analysis, the structural brain data (which is collapsed across time) must be optimally aligned. Ashburner (294) proposed a more elaborate registration process, which has advantages over traditional approaches. The DARTEL (Diffeomorphic Anatomical Registration Through Exponential Lie Algebra) algorithm for diffeomorphic image registration improves interparticipant registration and alignment of smaller central structures.

The preprocessed high-resolution T1 structural images are orientated to specifically place the anterior commissure at the origin of the standardised Montreal Neurologic Institute (MNI) co-ordinate system. Segmentation of the images into grey matter, white matter and CSF in native participant space is then performed, using the tissue probability maps. A rigid body transformation based on six parameters is then performed, with the DARTEL algorithm passing through six iterations of alternating between the creation of an average template of all the individual maps (i.e. a median image), and warping the individual participant images to the median image. This averaging and registration process simultaneously minimises a measure of difference between the image and the warped template, and the 'energy' measure (squared geodesic distance) of the deformations used to warp the template. The result of this process is to create a series of flow fields for each participant, which parameterise the deformation through non-linear registration of each individual image to the template for grey matter and white matter. The final average template can be registered to a standard anatomical template, enabling anatomical labelling.

Tensor based morphometry (TBM) can now be performed. TBM is an emerging computational analysis technique (295–299), which enables longitudinal quantitative assessment by identifying the region structural differences from the gradients of the deformation fields that nonlinearly warp each individual image to the template. TBM makes inferences based on Jacobian matrices of the deformation fields – these

values correspond to the relative volumes of tissue before and after warping. These matrices vigorously track local structural changes such as contraction (Jacobian determinant < 1) or expansion (Jacobian determinant > 1). The use of Jacobian determinants makes TBM advantageous as compared with voxel-based morphometry (VBM), particularly for small numbers of participants.

The data can then be transformed from participant 'median' space to standardised MNI space, with or without spatial smoothing. TBM analyses can be constrained using a ROI approach, as performed in Chapter 6, in order to increase the sensitivity of the statistical tests performed on the data.

2.8 Real-time fMRI methods

2.8.1 Overview

When designing a neurofeedback experiment using rt-fMRI, there are several factors that influence experimental design. These relate to the number of days of training, how long each training session should be, and over what period of time training should be distributed. Further, there are multiple options with regards to experimental set-up. It remains to be established what the optimum experimental designs are.

I have organised the methodological approaches I chose for both my rt-fMRI experiments into themes, including 'Computer factors', 'Physical factors', 'Brain factors' and 'Cognitive and behavioural factors'. I discuss the decisions made within the context of existing practices, and justify my choices in relation to alternative approaches.

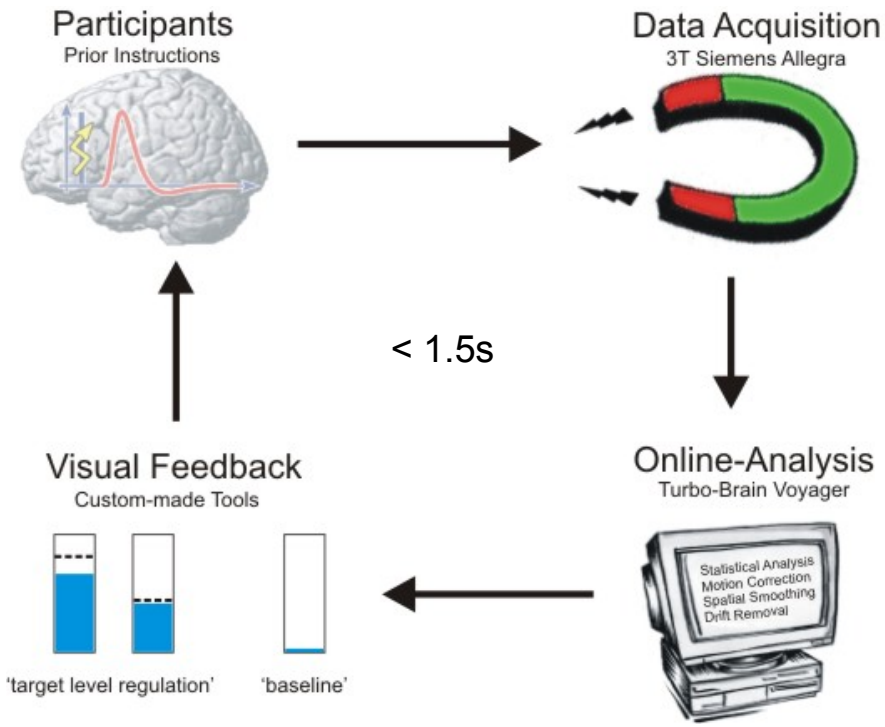

Fig 2-8. Real-time work-flow schematic. (See below)

There are four ubiquitous steps in a rt-fMRI BCI loop, which are:

1/ Participant instruction

2/ Imaging/data acquisition

3/ Online analysis

4/ Visual feedback

2.8.2 Computer factors

2.8.2.1 Setup and implementation

A popular software package for rt-fMRI experiments is the Turbo BrainVoyager (TBV) package in BrainVoyager QX v3.0 (Brain Innovation, Maastricht, The Netherlands) (300). I used this package in my experiments along with custom real-time image export tools programmed in ICE VA25 (Siemens Healthcare, Erlangen, Germany) (301), and custom MATLAB scripts (MathWorks, Natick, USA).

The real-time data preprocessing was performed with TBV and encompassed 3D motion correction with realignment to the preselected template, spatial smoothing, incremental linear detrending of the time series and SPM. The ROI timecourses were extracted from the prescribed ROI masks, then averaged and exported by TBV. Signal drift, spikes and high frequency noise were further removed in real time from the exported time courses with custom MATLAB scripts (302). The feedback signal was calculated and displayed to the participants in the form of a fluctuating thermometer bar. This display was produced using the Cogent toolbox (www.vislab.ucl.ac.uk/Cogent).

2.8.2.1 Delay of feedback

As discussed previously the haemodynamic signal related to neuronal activity occurs with a delay, necessarily introducing a delay in the BOLD signal relative to the cognitive processes that it reflects. Advances in computing power, and TBV

software, have reduced any additional processing delay associated with image acquisition, preprocessing and display, to 1 second, as was the case in my experimental set-up. (In neurofeedback studies, this delay has ranged from 2 minutes to a few seconds (40,232,301)).

Other options that may affect signal processing in real time include providing intermittent feedback, which may be presented at the end of a task block, or at serial points during the task block e.g. following the acquisition of several TRs, enabling averaging of the feedback signal. The suggested benefit of this is that more time points are available for filtering and signal averaging, and the participants therefore do no need to consider the haemodynamic delay. This approach also separates task strategy from interaction with, and evaluation of, the feedback signal, thus reducing distractor interference and increased cognitive load that this may cause. Although preliminary data (303,304) suggests that there maybe a more significant voluntary increase in brain activation with this method, it has not been consistently borne out (88,305). Current evidence supports the advantages of feedback being presented as soon as is possible after signal acquisition (72,85).

I opted to use a 'continuous, near real-time' presentation of the neurofeedback signal – I felt this to be more intuitive to the participants despite the presence of the haemodynamic delay. In the neurofeedback experiment presented in Chapter 6, all participants learned to control the signal, and did not report having difficulties with the associated haemodynamic delay.

2.8.2.2 Computation of feedback signal

In order to generate an online feedback signal, the time series must be sampled and analysed for which a number of approaches have been documented.

An *incremental GLM* was initially utilised (31). At any point during the scanning session, if an offline univariate analysis was computed on all the data collected up to that point, it would produce a similar result. Correlation analysis can be performed

cumulatively using the recursive least squares regression algorithm, and linear detrending can be performed incrementally by adding a linear predictor as a confound to the GLM (232). There have been subsequent optimisations (306,307) resulting in more rapid, and accurate online analyses, although the principle has remained the same. A further refinement of this technique has been instituted, with removal of high frequency noise, signal drift and spikes performed on the exported time series (302). I used this approach in my neurofeedback experiment (Chapter 6).

An alternative to the incremental GLM approach computes parameter estimates over a short temporal subsection of the data (308). Posse et al. (309) performed this *'sliding window'* approach to calculate the time series over 22 second blocks. The window progresses to keep pace with incoming scans, with the analysis performed after the completion of each block. The advantage of this method was an improved temporal resolution of the neurofeedback signal. I used this approach in the real time attention decoding experiment (Chapter 5).

2.8.2.3 Measures of the level of activity

A further consideration is how to use the timeseries data to produce the neurofeedback signal. In the majority of neurofeedback experiments, the participant is required to increase (or decrease) the level of activation related to a brain region, or a relative measure between brain regions e.g. an absolute difference or connectivity measures. The signal can be based on the mean percentage signal change (PSC) (58) as was the case in this thesis; other summary measures of activity can be used such as the median PSC. While these measures are comparable, the mean PSC is less affected by undesirable white noise, but can be susceptible to outlier voxel activation, which can occur in fMRI experiments. Hinds et al. (310) used a weighted average of the activity in modulated brain regions, explicitly modelling non-neural portions of the MR related signal (drift, white noise, head motion). They compared this against the aforementioned summary statistics including mean PSC. All measures were well correlated with the neurofeedback signal, but they found the best correlations with the actual neural signal, when using

a weighted average approach. The latter was specifically created using the voxels within the localiser defined ROI, but now taking into account voxel-specific efficiency estimates, by modelling the contribution of noise variance.

I used mean PSC as the dependent variable, and then calculated the difference in the levels of activation between two brain ROIs in the neurofeedback experiment (Chapter 6). In the online decoding experiment (Chapter 5) I examined the use of mean, mode, and median PSC as summary statistics. The mean PSC produced the most accurate online decoding of spatial attention deployment when I tested the procedure in an offline analysis of the data, providing the rationale for the use of mean PSC in both rt-fMRI experiments in this thesis.

2.8.2.4 Other measures of brain activation

Classifiers

Multivariate classification utilising spatially distributed activation patterns have also been applied in rt-fMRI paradigms. An initial timeseries is used as a training set, which is then compared to the task-related timeseries for similarity – the signal for neurofeedback is a measure of whole brain pattern comparison. A potential benefit of this approach is that it is more flexible. It does not need *a priori* hypotheses about functional localisation, enabling adaptive translations of brain activations in relation to individual strategies.

In Chapter 5, I describe an attention based rt- fMRI BCI, which decoded covert shifts of spatial attention on a single trial basis. In the preparation for this work, I applied multivariate classification approaches. However univariate decoding approaches were the most accurate, and therefore only these were used in the real time paradigm.

Connectivity measures

Although connectivity measures themselves can be used to produce the online neurofeedback training signal, I investigated *offline* DCM analyses to explore

network interactions underlying functional and behavioural changes following neurofeedback training in Chapter 6. In other studies, participants have been trained on the correlation coefficient between the BOLD time series in the two brain regions, to increase the functional connectivity between them (74). Kush et al. (12) implemented a neurofeedback approach utilising an online DCM package. Participants were trained to control the effective connectivity between ipsilateral primary visual cortex and superior parietal lobe, by using a covert attentional task to modulate a signal that was based in a Bayesian model comparison between two prior models of connectivity.

2.8.3 Physical/Scanning factors

2.8.3.1 Strength of the static magnetic field

The sensitivity of the BOLD signal and the signal-to-noise ratio can increase linearly with increasing field strengths i.e. according to the Boltzmann-distribution, magnetisation increases with field strength. Potentially, rt-fMRI paradigms may benefit from the increased signal-to-noise ratios at higher magnetic field strengths (19). However the gain may be offset by more field distortions and signal dropouts in the image due to larger frequency offsets in the object.

The rt-fMRI experiments in this thesis were all performed at 3T, which is currently the standard field strength for cognitive investigations. Rt-fMRI, however, has been performed successfully at 7T, in both decoding of brain signals, and neurofeedback control of brain activation in motor cortex (10,19,43–45,311).

2.8.3.2 Signal acquisition

For the rt-fMRI experiments in this thesis, whole brain functional images of the participant during task performance were recorded using EPIs, reconstructed and stored on the scanner console hard drive.

2.8.3.3 Echo time

The BOLD signal is dictated by the echo time, such that shortened echo times (as close to the T2* relaxation as possible) are optimal. Single RF shot multiple echo acquisition increases the BOLD contrast-to-noise ratio, while avoiding increased readout bandwidth (301). This multi-echo EPI sequence is optimal for rt-fMRI scanning as correction for image distortions is immediate, and not dependent on reference scans.

2.8.3.4 Image resolution

In terms of the resolution of an image, a compromise can be reached between reducing thermal noise by decreasing spatial resolution and increasing image read-out time, and keeping signal dropout and image distortions to a minimum. A typical fast fMRI protocol will use a 64 x 64 image matrix, with a 2-3mm in-plane resolution, as was used in the rt-fMRI paradigms in this thesis.

2.8.3.5 Head motion

Minor movements in the millimetre range can cause signal changes mimicking significant brain activation. Motion correction requires at least 3 slice acquisitions to reduce movement size to less than 1% of the voxel size in the case of a 64 x 64 matrix. Currently, ultra-fast motion detection camera technology is being developed to better detect motion for subsequent correction. This has particular value for realtime scanning of clinical populations e.g. movement disorders (Huntington's disease, (HD) (312) and translation of the experimental paradigms developed in this thesis.

2.8.3.6 Online spatial normalisation

This is an optional step in the TBV preprocessing package that I used, which typically consists of head motion correction, linear detrending, and spatial smoothing. As discussed earlier, it is used to address inter-individual differences, by warping

individual brains into a standardised anatomical space, utilising established templates such as the MNI template.

Standard offline spatial normalisation computations cannot be accomplished in the time frame of a single TR. Previous reported online techniques, have attempted to determine normalisation parameters offline following an initial preparatory run (313). A new advance in the technique, reported by Li et al. (243), has implemented a true online spatial normalisation method that can be accomplished with a short TR e.g. 2s. The method utilised an affine registration and a nonlinear registration based on cosine transforms.

2.8.3.7 Return of online data to the participant

The extracted data time series was exported via a customised visualisation interface to the participant inside the scanner. A delay is incurred during the course of the preceding two steps, and is ideally minimised to 1s (< 2s in my study). Note the data analysis and the feedback analysis together with visualisation are performed by two separate computers connected by a local area network.

2.8.4 Brain factors

2.8.4.1 Choice of brain region

Brain regions used most commonly in rt-fMRI studies have been those related to emotion processing e.g. amygdala (63,309,314), anterior cingulate (ACC) (64,232), posterior cingulate (PCC) (92) and anterior insula (315,316). Somatosensory regions are the next most commonly used, in healthy (58,65,303,317–319) and clinical populations (13,67). Sensory regions have been used and include auditory (66,320), and visual regions (12,19,17,16), and language (57) has also been tested. Higher order perceptual regions have, to date, only been used by Weiskopf et al. (301), and in my neurofeedback study (Chapter 6).

Brain activation utilised in the rt-fMRI workflow may be extracted in a number of ways. Regional brain activation can be used from an *a priori* determined ROI, usually identified with a functional localiser (85,321) as performed in both of my rt-fMRI experiments. Other options include connectivity measures and multivariate measures (see 'computer factors')

2.8.4.2 Predicting BCI control

Attempting to identify which participant volunteers are likely to achieve BCI control (i.e. voluntary control of brain activations) is important for establishing a successful study, but also has implications for clinical translation and patient selection. Although no predictive factors exist in the current rt-fMRI neurofeedback literature, some authors have attempted to predict BCI control in relation to anatomical structure. Halder et al. (78) found a positive correlation between individual SMR-BCI performance and deep white matter myelination. The mid-cingulate, and related white matter structure volume has been associated with EEG-BCI neurofeedback control of frontal theta activity. In clinical populations such as patients with amyotrophic lateral sclerosis, P300 attentional control was correlated with EEG-BCI

control. An important implication of predicting BCI performance would be targeting patients most likely to benefit from non-invasive BCI training, potentially serving as a prelude to the insertion of invasive BCIs. This could be combined with optimised neurofeedback training protocols.

2.8.4.3 Closed loop neurofeedback

An alternative neurofeedback approach is to covertly link a specific cognitive task with on-going brain activation in regions whose functional output is related to the task. By linking the task to the level of brain activation in a brain region (322), or to a whole brain pattern of brain activity using multivariate classification (323), a specific cognitive action can be trained. In both of these studies, category-specific attention (e.g. attending to house or faces) was trained by modulating a visual stimulus in relation to the strength of activation in the brain. This is termed closed-loop neurofeedback, and is a form of exogenous neurofeedback – the feedback signal is not directly presented to participant as in endogenous neurofeedback. Rather, it is used to trigger a linked stimulus which is not a direct representation of the neurofeedback signal, in order to manipulate brain activity and/or behaviour. In contrast, with endogenous neurofeedback as used in my own experimental neurofeedback work (Experiment 4 in Chapter 6) the participants are made explicitly aware that the visual interface (e.g. a fluctuating thermometer bar) is a direct measure of the neurofeedback signal.

2.8.5 Cognitive and behavioural factors

2.8.5.1 Role of mental imagery

Mental imagery is considered to be a critical component of neurofeedback learning. Brain regions related to visual mental imagery e.g. SPL, have recently been shown to undergo effective connectivity changes as measured with DCM, with trained brain

regions after neurofeedback (81). Whether or not a explicit instruction is provided to guide the participants, almost all studies documenting the manner in which participants gained control of the neurofeedback signal reported the use of visualisation and mental imagery techniques (232,301). Importantly, some well-controlled clinical and healthy volunteer studies have demonstrated a lack of a training effect when mental imagery alone is used, with or without follow-up practice (13,59), suggesting a qualitative difference between neurofeedback training that is mental imagery 'assisted' and mental imagery training alone. Practical examples of this difference have been provided with motor related mental imagery, as compared to actual movement. Overt hand movement leads to robust BOLD signal in motor cortex and supplementary motor cortex, whereas imagining motor activity frequently does not illicit any activation in some participants (324), and significantly less when its does (e.g. less than third of that related to overt movement (325)). In comparison, neurofeedback-led modulation of motor cortex activation, utilising mental imagery, was shown to achieve similar levels of activation to overt movement in motor regions (58). I provided participants with an explicit instruction to use of mental imagery in my neurofeedback experiment.

2.8.5.2 Neural underpinnings of the neurofeedback signal

The modulation of the neurofeedback signal has been primarily concerned with increasing or decreasing the overall level of BOLD activation within a target brain region, with a majority of studies focusing on increases in the level of activation (exceptions:- decrease/ increase (59), or decrease only (64,67,326)). A fundamental assumption is that, by increasing the overall level of BOLD activation in a specific task-related brain region, the underlying function of that brain region will change in a directional and predictable fashion. This assumption leads to a further series of rolled-up assumptions e.g. increased BOLD signals correspond to increased neuronal function, appropriate and non-conflicting neuronal populations have increased firing rates. Clearly this is likely to be a simplification – a particular brain region may have a variety of functions, only one or some of which are being targeted by neurofeedback training. A speculative explanation of the documented success

with neurofeedback in producing specific behavioral and perceptual changes may be suggested. Increased metabolic brain activity in the target brain region may serve to prime groups of neuronal representations related to a specific function of that brain region. By doing so, related task-based utilisation of these representations will occur in a more efficient and rapid manner, resulting in altered behavioural outputs.

It is possible that neurofeedback training of a specific regional BOLD signal results in learned control of the haemodynamic signal in a procedural manner e.g. control of capillary-arteriole-neuronal units (191) within a targeted brain region, to produce increased and co-ordinated neuronal activation, resulting in a more efficient, pruned neuronal responses and related functional output.

Control of the neurofeedback signal provides its own source of contingent reward. Although not performed in my own study, the neural activity underlying the neurofeedback signal can be more precisely 'shaped' online, either by providing an explicit reward e.g. monetary (88,327), or punishment, or by increasing the task difficulty associated with producing an increase in the level of the neurofeedback signal (323). Feedback can be provided during the task performance itself, helping to improve it. Andersson et al. (43,311) attempted online classification of covert shifts of spatial attention (right vs. left), with classification accuracy being fed back to the participants using changes in the colour of an arrow cue.

2.8.5.3 Showing the specificity of neurofeedback-based activation modulations

A fundamental requirement in neurofeedback studies is providing a means of confirming that the modulation of brain activity is specific to the targeted brain region, and is tightly correlated with the cognitive strategy employed to control the feedback signal. Effects that are not specific to the trained brain region/ strategy may arise from global increases in blood flow, or non-specific increases in arousal; this may result in anatomically widespread increases in the BOLD signal. Such changes can be mitigated by comparing the level of brain *activation* (cf. activity) in the modulated brain region against activation levels in another, carefully selected brain region. This

second brain region, a 'control' brain region (59,328), can be selected based on a the prior hypothesis is that it will not activate during the neurofeedback modulation task. The participant receives a feedback signal, which is based on the difference between the target region and the second region (59). An alternative to this is to utilise a second brain region that is activated by an orthogonal task. Here, participants can modulate the levels of activation in the two brain regions in a number of different ways, whilst still producing the same change in the direction of the neurofeedback signal[1]. This can result in a slightly complex biological interaction since the use of a differential signal can be manipulated in two directions i.e. bidirectional control. This can produce two mutually exclusive physiological responses, therefore providing a further internal control (329). This method has been used previously (57,58,315), and was utilised in the work presented in this thesis.

Further, the level of activation in the modulated brain region must be adjusted relative to a standardised baseline in order to measure *changes* in brain activation during neurofeedback. This can be done in one of two ways – it can be compared to a global mean, or it can be compared to activation levels in the target brain region during rest. The advantage of the former method is that it reduces the possibility of global brain activation changes during up/down regulation of the neurofeedback signal. As a consequence, I opted to compare the level of activation in trained brain regions with the global mean (Chapter 6).

2.8.5.4 Control group selection

An important consideration is confirming the effect of the neurofeedback training by selecting an appropriate control group. This is necessary, as the volitional engagement of the participants with a task may contribute to training effects, which

[1]Five potential activation states which could increase the difference between the two brain regions (R1 and R2) are possible, and will lead to up-regulation of the differential training signal (i.e. training signal = R1 - R2): 1) An increase in R1; 2) A decrease in R2; 3) a combination of 1) and 2); 4) a relatively greater increase in R1 as compared to R2; or 5) a relatively greater decrease in R2 as compared to R1.

are not specific to increased modulation of the neurofeedback signal. De Charms et al. (59) examined possible options for an appropriate control group. In a neurofeedback study on pain, four types of control groups were investigated. In a yoked control group, brain activation taken from participants in the experimental group was used as the neurofeedback signal. Although participants in the control group receiving this 'yoked' activation would be unable to gain control of the signal, it provided a constrained, and ordered dataset, potentially avoiding problems associated with failure of participants to engage with the experiment. This form of yoked control allows for the exclusion of effects related to expectation or suggestion.

In a second control group, the effects of attention were examined by providing instructions to overtly focus on, or away from an aversive stimulus. This group received no rt-fMRI feedback, engaging in behavioural training for twice as long as the experimental group. A third group learned to control activation levels in an alternative, unrelated brain region. Autonomic biofeedback was provided to a fourth group, with participants attempting to control skin conductance, heart rate and respiration. This group attempted to induce a more relaxed state by decreasing arousal, as compared to the experimental group who were trained to decrease the level of activation in a target brain region. In this study, only the group receiving feedback from the target brain region learned control of activation levels in this region, and demonstrated a reduction in pain perception. None of the control groups demonstrated a significant change in pain perception, with the yoked control group showing significant opposite effects. Although there are no other studies which provide an exhaustive exclusion of potential confounds, several studies have confirmed the lack of a training effect in relation to a yoked control signal, and mental imagery and behavioural training e.g. (13).

The use of an alternative brain region as described above, may provide the best neurobiological 'control', given that the two experimental groups can be matched for all variables, leaving only the neurofeedback signal as the differentiating variable. Thus, this approach has been utilised in previous studies (62,232) and was the approach I selected for the neurofeedback study in this thesis.

2.8.5.5 Selecting the number of neurofeedback training sessions

Control of the neurofeedback signal can be achieved after one day of training (2,301). The number of days used in rt-MRI neurofeedback paradigms typically varies from 1-5, with evidence of a training effect in all studies (Sitaram et al. (330): 3 days, Caria et al. (315): 1 day, Haller et al. (67): 4 days, De Charms et al. (59): 2-3 days). Within a training day, the number of sessions used has varied from 1-4, with a total in-scanner time of approximately 30-40 minutes. Feedback runs typically have task blocks of between 30-60 seconds, with interleaved rest blocks of 20-30 seconds, and are repeated between 3-5 times. Further considerations are the relative spacing of training days, and what, if any, behavioural training should take place between training days. Haller et al. (67) had 1 week between each of the 4 training days. Subramanian et al. (13) had 2 training days, with as much as 2-6 months between training sessions although during the interval, participants were required to practice the specific mental strategies that they engaged to control the neurofeedback signal.

It may be important for the consolidation of skill learning acquired during neurofeedback training sessions to be separated by a few days. Gulati et al. (331) showed that sleep, specifically slow-wave activity, helps to reinforce task-related neuronal ensemble activity after successful learning of neuroprosthetic skills. The authors of this study proposed that following neuroprosthetic skill learning, sleep helps in the offline processing of newly developed procedural skills and in the subsequent formation of a ‘neuroprosthetic procedural memory’. In my neurofeedback study (Chapter 6), the most significant changes in the control of the neurofeedback signal occurred between training days, potentially supporting a role for sleep. Although there has been no strong proposal for a benefit or lack thereof in multiple rt-fMRI neurofeedback training days (although see (332), a suggestion can be made based on my own study i.e. 3-4 training sessions, containing 3-4 training blocks, and lasting no more than 45 minutes, and spread over the course of 3-4 days. This may provide an increased likelihood of a significant training effect, while remaining practical in terms of scanning time and what can be demanded/expected of participants and patients.

2.8.5.6 The training related 'learning' curve

The term relates to the graphical representation of learning effects with increasing training experience, first coined in 1885 by Herman Ebbinghaus (333). In rt-fMRI neurofeedback, this provides an indication of the levels of brain activation as a function of the amount of neurofeedback training e.g. mean PSC in a trained ROI, plotted across training runs or sessions i.e. each time the scanner is switched on/off. It therefore can provide an indication of how successfully a brain region was trained with neurofeedback. While intuitively a linear trend would be expected as training increases, a number of possible patterns may be observed e.g. cuboidal, which may be influenced by the occurrence of breaks, or the effects of intervening sleep e.g. consolidation vs. forgetting. For the work presented in this thesis, I plotted the mean PSC of the neurofeedback training signal across the 9 training runs, to examine for a training effect.

2.8.5.7 Use of instruction

This is potentially the most important component of a neurofeedback study (72,90,334). The provision of instructions on signal control strategy have been shown to produce steeper learning curves (62), as well as enabling learning in participants who were previously showing no signs of voluntary self-regulation. This has been supported by a head-to-head comparison between participants learning to control language areas. Subjects who received instructions on specific mental imagery strategies learned voluntary modulation of brain activity, whereas participants who received no instruction were unable to learn at all (335). Scharnowski et al. (62) confirmed participants gained control of the neurofeedback signal more rapidly when provided with an explicit instruction linking the functionality of the brain region being modulated to the cognitive strategy employed, resulting in an in-experiment change of methodology. However, Caria et al. (315) have demonstrated that participants can learn voluntary self-regulation in the absence of explicit instructions, within one day of training. Previous healthy volunteer studies (186, 220) and recent rodent work (94) confirmed learnt self-regulation of brain activation as a function of operant conditioning, with a lack of explicit guidance.

Therefore, control of the neurofeedback signal in the absence of instruction linking metal imagery with the function of the brain region to be trained is possible. Nonetheless from an operational point of view, accelerated and more robust learning of voluntary modulation of the neurofeedback signal is aided by explicit instruction. This formed the basis of my reasoning to include this as part of the methodology in my neurofeedback experiment.

2.8.5.8 Feedback signal settings – modality, parametric, delay

The substrate of the feedback signal can be any sensory modality e.g. auditory, tactile, or a combination e.g. auditory interfaces (309,337). EEG feedback studies have indicated that visual feedback is the most successful e.g. (338). A variety of visual interfaces can be used, such as simple scrolling graphs (232,301), 'mercury' thermometers (58,315,316,16,339), circles of changing diameter (17,340), and numerical values (304), to more representative displays such as fires of differing sizes (59), and emotional faces for social reinforcement (89). With regards to the latter class, virtual reality interfaces have also been attempted (341–344). Most recently de Bettencourt et al. (323) utilised the stimulus that drove the target brain area as the display itself, strengthening either face or house stimuli in an overlapping display depending on the type of attentional feedback control required. Finally, visual feedback is considered to be superior to auditory feedback (345), and should be presented parametrically, rather than in a binary form (346), the former of which was selected for my neurofeedback experiment.

2.8.5.9 Modulating brain activation after neurofeedback training

Following neurofeedback training, it is often desirable to show that participants can continue to control the level of brain activation in the trained brain region/regions without concomitant neurofeedback. On the other hand, participants may only be able to increase the activation levels in the presence of a feedback signal

(155,165,185,186). The ability modulate without feedback can be tested with the use of a transfer task, or by asking participants to attempt to concurrently ‘up-regulate’ brain activation while performing a behavioural task.

A transfer task can assess a central question in the brain training literature as to whether neurofeedback training can fundamentally alter the output of a trained brain region after training (349). It involves the application of a task, which can be motor or cognitive, to test related but not identical functional domains of the trained brain region/s. It is also possible that the transfer task is performed while the participant modulates brain activity as learned during neurofeedback, but now in the absence of a neurofeedback signal (16). This was the approach I used, in which a transfer task was administered, with and without concurrent modulation of brain activity following neurofeedback training.

Successful post-training transfer should be evident as an individual’s ability to modulate the trained brain region on the transfer task. It should, however, be considered that the absence of a transfer effect may be as a result of methodology or training regime (e.g. single training sessions vs. multiple sessions) or related to the neurobiology of the specific brain region being trained.

ABSTRACT: Experiment 1

Multivariate classification of spatial attention in higher order visual areas

Introduction

I explored classification of brain activations associated with the deployment of covert attention to cued spatial locations. Region-specific activations were used including primary retinotopic cortex, and high order visual areas i.e. parietal cortex, FEF, and lateral occipital cortex (LOC). I tested the hypothesis that this information could be 'decoded' using multivariate approaches, with accuracies that would be comparable to classification with standard univariate techniques.

Methods

Thirteen neurologically normal adult volunteers underwent fMRI scanning during which they were cued to attend stimuli were presented simultaneously in each of the four visual quadrants. Central eye fixation was required and monitored with an eyetracker. BOLD responses to the task versus rest were used to generate functional ROIs representing retinotopic and higher order visual areas responsive to the stimuli. These ROIs included primary visual cortex, parietal cortex, FEF, SMA and LOC. For the attentional conditions, I used multivariate classifiers (Correlational classifier (CC), Linear Discriminant Analysis (LDA), with/ without principal components analysis (PCA), on the spatial patterns of activation produced within each ROI. This generated classification accuracies for the deployment of spatial attention for each of the conditions i.e. 'four directions', 'left versus right', 'up versus down', 'diagonals'.

Results

The three multivariate classifiers performed significantly above chance in distinguishing 'left versus right', and for four quadrants attention; this was most robust in primary visual areas e.g. bilateral 'occipital': 4 directions classification accuracy; **CC**: 30% $p<0.01$ $t(12)=4.74$, **LDA**: 31% $p=0.01$ $t(12)=2,78$, **LDA with**

PCA: 30% p<0.01, t(12)= 3.12, and in higher visual areas i.e. bilateral 'LOC' **CC**: 31% p<0.01 t (12)= 3.90, **LDA**: 30% p=0.04 t(12)=1.96, **LDA with PCA**: 28% p=0.02, t(12)= 2.39. Bilateral parietal, **CC**: 29% p=0.02 t (12)=2.44, **LDA with PCA**: 27% p=0.08, t(12)= 1.53. Examining 'diagonals' and 'up versus down' conditions provided above chance classification for both only with the bilateral Occipital ROIs using CC, and LDA with PCA classifiers. Both CC and LDA with PCA classifiers performed best, as compared to LDA, using the univariate approach, no significant classification accuracies were achieved in any studied regions.

Conclusions

I was able to apply multivariate analyses to classify the direction of spatial attention deployment to stimuli in four quadrants using task-related spatial patterns of BOLD activation in visual regions including high order visual areas. I provide evidence for the presence of spatial information in the LOC. This work provides the basis for establishing a 'decoding' algorithm for the deployment of spatial attention, with a view to developing practical applications such as an 'attention driven' BCIs. The purpose of this would be to provide a communication platform for patients who may be unable to speak and/or move.

Key words:

rt-fMRI, BCI, covert spatial attention, classification, decoding,

3. EXPERIMENT 1

Multivariate classification of spatial attention in higher order visual areas

3.1 Introduction

In the stimulus-rich visual world, our brains must ensure that limited cognitive resources are appropriately allocated. One of the processes for doing this is visual attention. The importance of visual attention in the processing of relevant perceptual input has been well documented. For example, psychophysical experiments on participants with change blindness, requiring them to identify changes in a scene in the absence of a cue, have reported the difficulty participants experience in identifying change, unless attention is directed to the locus of change. In terms of the neurological underpinnings of attention, patients with strokes affecting the fronto-parietal cortex with damage to putative components of the dorsal circuit of attention i.e. FEF and IPS (350–352), show significant deficits in spatial attention. Lesions specifically of parietal regions demonstrate profound supramodal attentional deficits as a result of both contralateral and ipsilateral lesions (353–356).

Attention may be saccade-mediated i.e. overt attention, or may take advantage of the peripheral representation of the visual field on the retina i.e. covert attention (357). Both processes may be voluntary or involuntary (358). Examining the voluntary deployment of attention more specifically, it relies on intrinsically driven top down processing (i.e. 'endogenous attention' as compared to 'exogenous attention' which is driven by an external stimulus). The mechanisms governing 'top-down attention' have implicated the possibility of a putative 'attention command signal' , which may direct or drive the entire process. Identification of such a neural command signal, if it exists, has practical value. It provides direct, and potentially time-critical information on the manner in which an organism interacts with its environment. Extracting such an attention command signal, might enable it to be

used as surrogate for explicit verbal or motor communication, potentially paving the way towards targeted therapeutic interventions. This could, for example, be an interactive communication prosthesis that relies on manipulations of visual attention (15). Such a prosthesis might use a classification algorithm to read targeted brain activation, serving to ‘decode’ the direction of the attention command signal as covertly directed by the user. This translated signal could be used to navigate a visual interface through a selection of choices linked to real-world actions i.e. a speller(39,340,359), images of food or drink objects. This device might additionally help in the clinical domain, by bypassing brain injuries, downstream cortical processing, and disabled or damaged target organs e.g. paralysed limbs, dysfunctional speech.

In this chapter, I seek to establish the basis for an attention-driven communications interface. In setting out to do this, an understanding of the neural circuitry of attention is indicated, together with examination of its interaction with the subject of attentional deployment i.e. spatial attention, object based attention. I introduce key findings from the attention literature below, and then outline analysis methods which may be best suited to optimise the identification and use of neural information for the purposes of an fMRI based, attention-driven communication interface.

An existing conceptualisation of attentional circuitry separates it into dorsal and ventral streams, with the fronto-parietal network acting as a ‘source’ of top-down signals exerting influence over the ’sink’ of the occipito-inferotemporal cortex (360). Lesion studies in patients have provided support for the anatomical segregation of these systems. Right parietal damage has been associated with impairments in the endogenous orienting of contralateral spatial attention, with a relative sparing of exogenous orienting (361); in contrast, damage to the temporo-parietal junction results in extinction-like patterns, which are related to exogenous orienting. Corbetta and Shulman (105) put forward a model of spatial attention, where goal-driven control is reliant on the activity of a bilateral fronto-parietal network, whereas a ventral fronto-parietal network is responsible for salient attentional capture. (NB: There is, however, recent work by Vossel and Geng (362), offering an alternative to this model).

The deployment of visual attention implies the presence of representations of visual space in those cortical areas under the influence of visual attention. The neuronal organisation of primary sensory cortex reflects the close spatial proximity of neurons that are involved in the same sensory computations. Considering the domain of visual attention, primary visual cortex (including V1, V2, V3, V4 (284),V3a (363), V3b (364)) has retinotopic organisation, based on neurons with small well defined receptive fields. Higher order visual areas tend to be smaller in size, and have neurons with larger receptive fields (365,366). The amplification effect of attention increases along the visual hierarchy i.e. V1 to V4 and beyond, with increasing signals associated with sensory processing, to create an attentional gradient. This effect is greater in higher visual areas than in primary visual cortex, potentially because the widening summation across progressively larger visual fields in higher areas, results in increasingly steeper contrast response functions (367). Spatially specific modulations of visual attention have been identified in a number of higher order visual areas - in particular those regions in the dorsal fronto-parietal network classically considered to serve as command centres in the top down organisational hierarchy. These include the FEF, precentral gyrus, inferior frontal sulcus, IPS (differentiated into 6 subunits from IPS0 to IPS5) and the SPL (112,368–378).

Higher order visual brain regions such as the LOC and inferior temporal cortex, have shown evidence of multiple retinotopically organised areas, giving rise to the possibility of their role in spatial attention, linked to their roles in object processing(379). These areas are part of the ventral ‘what’ pathway, and have general purpose (e.g. object selectivity in LOC (380,381)) and category-specific roles (e.g. face selectivity in the FFA (382)) in object processing, in the absence of attention. The LOC was first described functionally in 1995, located on the posterolateral aspect of the fusiform gyrus, extending ventrally and dorsally (380). Object selectivity was demonstrated in this region, which was subsequently confirmed as being independent of low level visual cues and specific stimulus categories (381). Retinotopy in the LOC was demonstrated using phase-encoded mapping techniques in functionally defined ROIs(112). A more direct illustration of retinotopy in the LOC was provided by the application of electrical cortical stimulation in epilepsy patients with subdural electrodes- the authors elicited phosphenes in patients in a pattern demonstrating retinotopic representation (383). Recent additions

to this literature have delineated both central (foveal) and peripheral representations of objects in the LOC such that while position invariance may occur (48,384) there is evidence for a co-ordinate system which is retinotopically organised, and may provide differential information about where a particular object is in space. The emergent neurobiological question which arises, given the presence of retinotopic circuitry in LOC, is whether 'attentionotopy' and the application of spatial attention can be identified in this region. The concept of 'attentionotopy' relates to topographic maps of spatial attention in higher order visual regions, based on the presence of retinotopic maps in these regions. While retinotopic maps have been demonstrated throughout the visual pathway, early visual areas show predominantly stimulus – driven retinotopy, higher order visual in frontal and parietal areas demonstrate predominantly attention-driven retinotopy (38,369). These higher order areas may send spatially-specific top-down attentional command signals which act on early visual areas. In addition to higher-order brain regions known take part in the neural action of spatial attention, I specifically investigate the presence of neural signals related to the deployment of spatial attention in the LOC in this experiment. Bilateral ROIs are explored for all brain regions, as retinotopy in both dorsal and ventral visual areas have been consistently shown to be present bilaterally, and/or in both hemispheres with no clear lateralisation (41,374,379).

The accuracy with which spatial information present in higher order visual areas such as LOC, can be identified from fMRI imaging poses a challenge given the spatial resolution of MRI. The BOLD signal of an imaged voxel (in the case of the studies in this thesis, $3mm^3$) reflects the averaged activity of several thousands of neurones (385) which may reflect heterogeneous neural activity in relation to the cognitive task and behaviour under study. This 'blurring' of neural information is increased further by the application of traditional analysis approaches that spatially average across only those voxels that contain responses that are statistically significant with respect to the experimental task. While this has the advantage of reducing 'noise', it also reduces 'signal' by excluding voxels with weaker (i.e. non-significant) but potentially contributory information, and ignoring fine-grained spatial patterns that exist across voxels (264). This may be of relevance when examining neuronal populations that encode topographical space in a contiguous manner. The use of multivariate analysis on a voxel-wise basis seeks to avoid these concerns by

using pattern classification techniques to examine the spatial patterns of responses taking place across multiple voxels, including those which have non-statistically significant responses to the task, when measured using univariate methods. I specifically use multivariate analysis of the imaging data in this experiment, and compare it with traditional univariate analysis techniques, in order identify which is the optimal means of confirming neural activity linked to the covert deployment of spatial attention in dorsal and ventral higher order visual regions.

MVPA employs machine-learning techniques, including classification algorithms, to examine spatially distributed patterns of brain activations. The use of classifiers aims at establishing above-chance 'decoding' or 'classification accuracies' - that is reliable patterns of activations that contain information that can be used to determine the specific experimental condition. There are, however, several types of classifier, each with differences in how the data is treated, and which can therefore produce different levels of classification accuracy relative to one another. All classifiers have in common the use of a training data set to define a decision boundary in the space of response patterns e.g. the space spanned by the activity levels of the voxels in the ROI (269). Classifiers can differ in the shapes they allow for the decision boundary (e.g. hyperplanes in linear classifiers; more complex non-planar boundaries in nonlinear classifiers) and in the way the boundary is placed on the basis of the training data. For example, LDA places the decision hyperplane so as to optimally discriminate two equal-covariance Gaussians; a linear SVM places a decision hyperplane so as to maximise the margin to the patterns on either side. Each classifier has certain strengths and weaknesses along this axis of interest (e.g. more flexible decision boundaries lead to better separation of data, but introduce the risk of over fitting (266)). Further, each classifier must be considered with respect to the brain region being studied, the use of ROIs, single voxel response estimates (beta estimates or t-values) and categorical dichotomy. There is recent empirical evidence to suggest that linear classifiers perform better than non-linear classifiers, with some preference for LDA and pattern correlational classifiers on account of simplicity, robustness, interpretability and ease of application (269).

MVPA together with machine-learning algorithms have been widely used to examine the representational content in patterns of neural activity (38,256–258,261–264).

However it is currently unknown as to what approach might be best suited to classification of a particular cognitive behaviour of interest that relies on specific ROIs. This is particularly important for the design and implementation of a neural communication device in which accurate, robust 'decoding' from selected brain areas is critical for the success of communication. The aim of this study was therefore to examine whether the deployment of spatial attention can be classified or 'decoded' using multi-voxel pattern classification techniques, at above chance level in higher order visual areas.

In this experiment, I examined brain activations in higher order visual areas including LOC, associated with the deployment of covert, 4 quadrant spatial attention to cued locations. I was specifically interested in testing the hypothesis that region-specific activations could provide information with regards to the deployment of spatial attention, and that this information could be successfully extracted by using univariate and multivariate classification techniques. I predicted that across the 3 different classifier approaches tested, multivariate approaches with prior feature reduction (e.g. LDA with PCA) would decode with the highest levels of classification accuracy, the location of an individual's covertly deployed spatial attention, in higher order visual areas.

3.2 Methods

3.2.1 Participants

Thirteen neurologically normal adult volunteers (24–32 years of age; mean age = 28; 8 females) with normal or corrected-to-normal visual acuity were recruited from the general population to participate in the experiment. Each participant was provided with written informed consent approved by the local ethics committee (UCL Ethics Committee approval number: 09/H0716/14) and passed a MRI safety medical

screening approved by the Wellcome Trust Centre for Neuroimaging (where MRI scanning was performed).

3.2.2 Experimental set-up

Stimuli were back projected (NEC Lt Projector) onto a screen mounted to the top of the magnet bore behind the participant's head. All visual stimuli described below were generated and displayed via MATLAB using the Cogent toolbox (http://www.vislab.ucl.ac.uk/cogent_2000). Participants viewed the screen (via a mirror mounted to the head coil) at an optical distance of 52cm. The participants responded via a pair of custom-built, MR-compatible, fiber-optic push button response boxes.

3.2.3 Experimental procedure

The experiment was divided into 12 runs. During the course of one run, there was a cue presentation (1s) followed by 16 presentations of the stimuli (0.5s each), each followed by a response interval (1.5s) and an interstimulus interval of 1.5s. Each block of 'cued' presentations (i.e. 16) was followed by a rest interval (16s) (See Fig 3-1). Each run included cues to all 4 quadrants in the visual field (upper and lower right and left quadrants) displayed in a pseudo-randomised order.

The stimuli consisted of 4 black-and-white spirals of identical size (diameter 2 degrees of visual angle (VA)) and set at identical eccentricities (6 degrees VA) from the centre, in 4 separate quadrants of the display. The spirals differed in their chirality. During the course of a cued run, all 4 quadrants displayed spirals, with the participant attending the spiral (and its associated chirality) in the cued quadrant only. To ensure attention was maintained, the participants were asked to indicate on each trial whether the attended spiral was right or left handed through a two-button box press.

The sequence of the chirality of the spirals was pseudo-randomised through the use of a balanced Latin square. The sequence of the attentional cues across the 12 runs was pseudo-randomised such that each cue was used four times in a non-sequential fashion.

3.2.4 Eye tracking

Throughout the experiment the participants were instructed to maintain central eye fixation at all times, centred on an oval dot at the centre of the screen. This was present throughout the task and the rest interval. Central fixation was ensured through the use of eye tracking, with eye movements being recorded for all thirteen participants, throughout each scanning run. Eye position and pupil diameter were sampled at 60 Hz using long-range infrared video-oculography (ASL 504LRO Eye Tracking System, MA, USA).

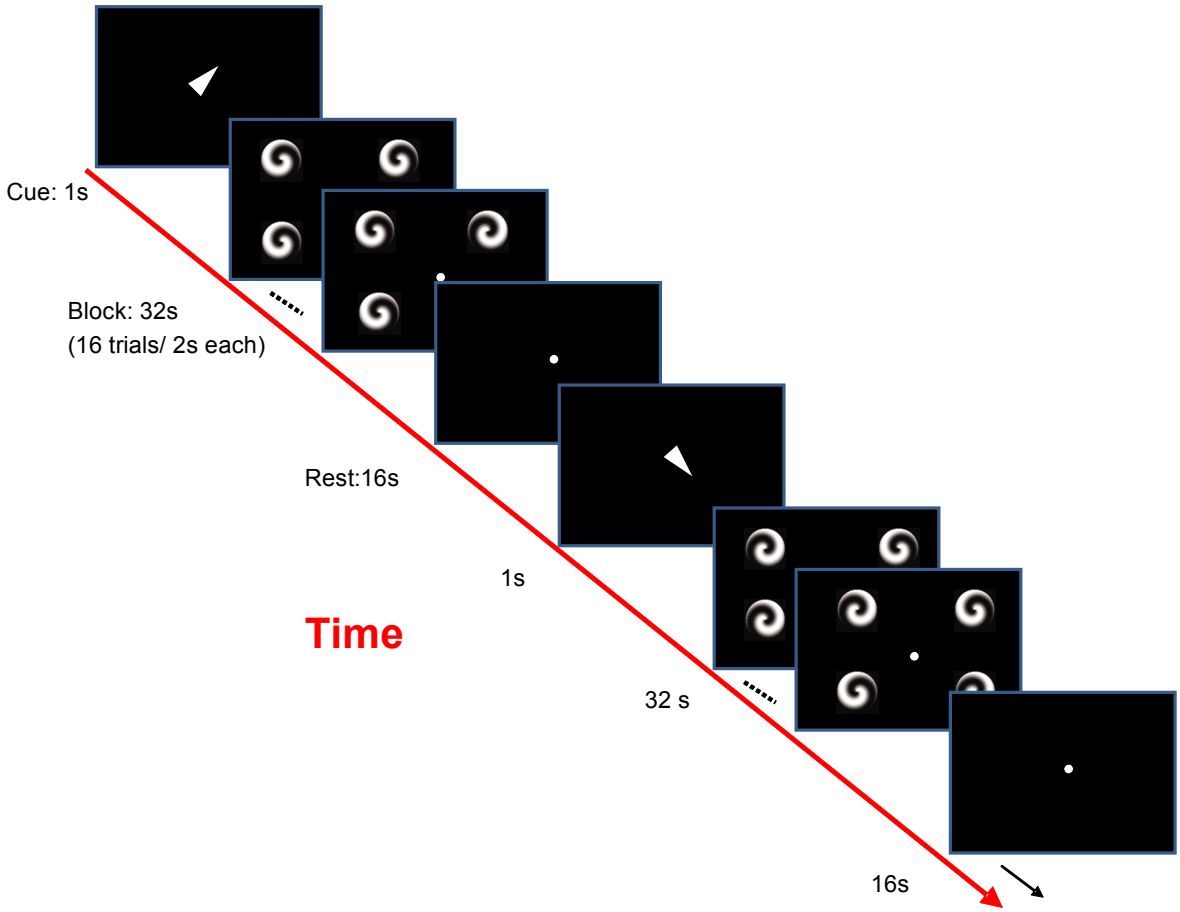

Fig 3-1. Experiment 1 task presentation, as visualised in scanner display.

A schematic of the task requiring participants to attend to a quadrant, cued at the start of a block, followed by 16 trials in which the chirality of a spiral in the cued quadrant only was indicted by a button-press response. The block of trials were followed by a 16s rest block, at the end of which the spatial cue was re-presented, indicating attention to another quadrant.

3.2.5 Data acquisition

FMRI data were acquired on a 3T Allegra scanner (Siemens, Erlangen, Germany). Structural images were acquired using a T1-weighted sequence (in the sagittal plane, 176 partitions, TR= 7.92s, TE=2.4ms, FOV=256x240). Functional images were acquired with a gradient-echo EPI sequence (TR= 60ms, TE=30 ms, FOV=64x64, matrix=64×72, interleaved acquisition, no gap, 3mm isotropic resolution, 32 slices). Slices were positioned along the slope of the temporal lobe and covered ventral visual cortex and parietal cortex. The sessions of the experiment consisted of 205 volumes each. Fieldmaps were acquired to correct for geometric distortions in the EPI images related to inhomogeneities in the magnetic field.

3.2.6 fMRI Data analysis

3.2.6.1 Basic Pre-processing

To allow for T1 equilibration, the first five images of each run were discarded. Pre-processing of the data involved mean bias correction, realignment (of each scan to the first scan of each run), unwarping, and co-registration of the functional data to the structural scan. Data were smoothed with a 6mm FWHM Gaussian kernel. The data were filtered with a standard 128s cut-off, high-pass filter to remove low-frequency noise including differences between runs, while at the same time preserving as many of the spontaneous fMRI fluctuations as possible (386). I used session specific grand mean scaling, with no global normalisation when preparing the GLM in SPM 8 (Wellcome Trust Centre for Neuroimaging, UK). Movement parameters in the three directions of motion and three degrees of rotation were included in the GLM as confound variables.

3.2.6.2 Defining functional ROIs

For each participant, the parameter estimates resulting from the fit of the GLM to the fMRI time series were used to calculate a t-statistic at each voxel indicating a statistical probability of task-related activation. This resulted in participant-specific statistical maps of t-values summarising activations associated to the task (in this case summed across all conditions) compared to rest. The t-maps were overlaid onto each inflated cortical hemisphere (reconstructed from the anatomical images) using Freesurfer (http://surfer.nmr.mgh.harvard.edu/fswiki) (285,292). Specific ROIs were then defined and included the parietal cortex (superior and inferior), FEF, polar occipital cortex and LOC.

3.2.7 Data analysis- Eye tracking

Eye tracking data were analysed with in-house scripts utilising MATLAB and SPM8. After eye blinks and periods of signal loss were removed from the data, mean and median eye position, expressed as a distance from central fixation were calculated. This was performed for each trial type and for all thirteen participants. A repeated measures analysis of variance (ANOVA) was used calculate whether mean eye position deviated significantly from fixation, for and between conditions.

3.2.8 Multivariate analysis

3.2.8.1 Preprocessing

FMRI data were spatially smoothed (6 mm kernel FWHM). As has been outlined by a number of authors, the spatial pattern information exploited by many multivariate pattern classification analyses is represented on a relatively coarse spatial scale

(387–390). Consequently, spatial smoothing at an appropriate level improves classification performance (391).

I modelled each of the attentional conditions, together with regressors for ‘cue’ and for ‘rest’ (across the course of a whole run). The resulting parameter estimates for each condition, and every run were used to generate t-maps, created using contrasts of each condition against the global mean. The voxels belonging to each ROI were identified by projecting the labelled surface vertices back into voxel space. For each ROI, the t-values for each voxel were extracted and made into a vector.

3.2.8.2 Overview

In a manner similar to the application of regression when setting up a GLM in SPM, in order to establish the time series associated with each voxel (from many columns in the design matrix), classifiers work in the opposite direction, predicting parts of the design matrix from groups of voxels. A particular classifier will take various features (independent variables or predictors in regression; in this experiment it would be voxels) in an ‘example’ (the set of independent variables), and predict the ‘class’ that example belongs to (the dependent variable i.e. the stimulus used to generate task-based BOLD activations). The classifier will have a set of parameters that are ‘learned’ from the training data effectively building a model defining the relationship between the features and the class ‘label’ in the training set. The trained classifier can then determine whether the voxels being tested (e.g. from an ROI) have information about the class of the example; the classifier is used to the test the feature/class relationship by applying its function to different sets of examples, the ‘test’ data. The latter assumes that the training and test examples are being independently drawn from the example distribution.

Prior to training the classifiers, the fMRI data must be transformed into examples. There are a number of ways of doing this, ranging from using the activity signal itself from one or more voxels, to utilising the linear model activity estimate from the GLM to predict each voxel- so called ‘single voxel responses’. These may be beta or t-

values for a given condition. As outlined in the pre-processing steps, I opted to use t-values, as previously supported by Kriegeskorte (269).

3.2.8.3 Feature selection

Given the typically large number of features in relation to the number of examples in a fMRI study, feature selection or feature reduction is often applied to the data set. This prevents the phenomenon of 'overfitting'- essentially a multiple comparisons problem, in that it may be possible to find a classification function that accurately classifies the examples in the training set, without it being effective when applied to the test data. Simple linear classifiers also have less tendency to over-fit the data which is the rationale behind the use of correlational and LDA classifiers in this experiment. This has been supported more formally in the literature (265,266,269). I addressed feature selection in a number of ways :-

1. Regions of Interest

Rather than considering all the voxels in the brain, I limited the search to within a specific ROI. Furthermore, I used functional ROIs rather than anatomical ROIs, restricting analysis to those voxels that demonstrated above-threshold activity during task trials relative to the baseline (modelled as rest). I defined ROIs from within high order and primary visual cortex, including frontal, parietal, lateral occipital and occipital regions. Bilateral ROIs were created, with all conditions being decoded using these ROIs. I did not create ROIs based on voxels that specifically distinguished an attentional condition i.e. to one particular quadrant, although this has been done within the context of operational brain-machine interfaces (19,43,44,311,340).

2. Dimensionality Reduction

This involves reducing the original feature space into a new, lower dimensional feature space, yielding a dataset matrix with the same number of rows but a reduced number of columns. I used a PCA/ singular value decomposition (392) with an LDA

classifier, selecting eigenvectors that explained more than 5% of the variance in the data.

3. Ranking

An application of the 'scoring/filtering' method involved ranking each of the features by a given criterion- each feature was scored by itself in an analogous way to univariate tests of a voxel. The 'best' voxels were selected- I picked the best 60 voxels for the multivariate classifiers. In the present study, ranking was performed on the data in the training set, and then applied to the data in both the test and training set. An F-test was used to compare the two conditions being classified for each voxel (e.g. left vs. right etc.); following this all voxels were ranked in descending order in relation to their F-statistic. This would then be dependent on the number of voxels from a particular ROI that were used in the classification procedure.

3.2.8.4 Cross-validation

To establish a useful estimate of the classifier's accuracy, the optimal situation is to train and test on as many examples as are available (assuming ideally that the classes are balanced, and that the training data in each fold contains examples of all classes (266). I used a 'leave-one-out' cross-validation (268), effectively producing as many classifiers as there were examples.

3.2.8.5 Experimental Analysis

Multivariate classification was established by using a leave-one-run-out cross-validation procedure in which samples from all except one session were assigned to a training set and the remaining samples were used as a test set. This step was repeated across all possible run combinations, such that each run was used as the test set. Classification performance for each cross-validation was estimated as the percentage of correct classifications, and the final classification accuracy was calculated by averaging performances from all cross-validations.

I classified, in separate analyses, the following attentional conditions: - left vs. right, up vs. down diagonals (i.e. top-left plus bottom-right vs. top-right plus bottom-left) and four-way. For the binary classifications I disregarded the irrelevant dimensions i.e. when classifying left vs. right, upper and lower visual field attention was disregarded.

I determined a priori the number of voxels to be considered in the classification step. All results presented were based on the use of 60 voxels (considered to be broadly representative of the ROI). Previous authors have demonstrated that fewer, but visually more responsive voxels, lead to improved classification accuracy (269,393). Although accuracies degraded at larger numbers of voxels, this was also the case with fewer voxels i.e. under 30. There is likely to be a threshold related to the volumetric size of the anatomic ROI, and its task-related functionally defined analogue.

3.2.9 Choice of classification technique

3.2.9.1 Multivariate approaches:

All classifiers have in common the defining of a decision boundary (hyperplane) in the space of response patterns (that is the space spanned by the activity levels of the voxels in the ROI). I used simple linear classifiers with relatively rigid hyperplanes, for ease of application, comparison and avoidance of over-fitting.

The specific classifiers used were a correlational classifier, LDA, and LDA with PCA. The use of PCA with the LDA classifier was because there were not enough examples to estimate its covariance matrix reliability (270).

1. Correlational Classifier

For each attentional condition I calculated the mean across all runs in the training set. To decode I calculated a linear correlation between each sample in the test set and the mean samples from the training set. A test sample was then assigned to the condition that produced the greater correlation coefficient.

2. Linear Discriminant Analysis

This classifier determines the discriminant dimension response pattern space, in which the ratio of between-class over within-class variance of the data is maximised (269). After projection of the data into the linear discriminant dimension, a classification threshold is placed at the midpoint between the two class means. The classifier is Bayes-optimal in that it ignores estimation error, and assumes Gaussian within-class distributions.

In a two class-classification problem as within our experiment, the normal vector of the decision boundary is estimated as a product of the within-class covariance matrix and the mean vectors of each class. The side of the hyperplane that a discriminant value falls on determines which of the two classes it belongs to (after establishing its distance from the hyperplane in relation to the mean value of all samples, and the vector of voxel values, whether it is a positive or negative integer).

The complementary aspect to this method in our analysis was the use of PCA for dimensionality reduction, producing eigenvectors (from the individual voxels) as the separable features.

3.2.9.2 Single Feature analysis:

1. Correlational Classifier with Principal Components Analysis (single most discriminant eigenvector)

Following feature reduction using PCA, correlational analysis was then performed using the single most discriminating eigenvector when considering the two conditions of interest.

2. Univariate approaches

In order to contrast and explore the strength of the multivariate classification I examined classification based on single features, representing a special case of the classification methods outlined above, but now with extreme feature reduction. The motivation behind the multivariate approaches was a finer grained modelling of the information provided by the experimental data set. Nonetheless I sought to examine whether a broader approximation of the task related activations would still provide similar levels of classification accuracy.

i) Average of t-values across a ROI

The univariate analysis was performed using the average of the t-values across each ROI. I established classification of the average signal in an ROI in relation to the condition under examination. Once again ranking was used to isolate the 60 most informative voxels with regards to the two conditions examined, with the average calculated across the t-values from these voxels only.

ii) Correlational Classifier with single 'best' voxel

This identified the single most discriminative voxel in each ROI following ranking i.e. the voxel that had the largest F-value when comparing the two conditions of interest in the training set.

3.3 Results

3.3.1 Overview

I used multivariate classifiers as well as routine univariate classification to determine whether activity within each ROI could be used to classify the direction of spatial attention in each of the four attention conditions and for combinations of the conditions i.e. left vs. right, up vs. down, and diagonals. Chance level of performance was therefore 25% for the four-way classification and 50% for the paired attentional comparisons.

Potentially the most important comparisons for a real world setting, e.g. navigating a room, were left versus right and the four-way attention conditions (offering the largest number of options).

3.3.7 Eye position data

Participants were requested to maintain central eye fixation throughout all of the 4 task blocks, for each of the 12 sessions. The repeated measures demonstrated no statistical difference in mean eye position from fixation or between conditions. An ANOVA was performed for x- and y-axis for the thirteen participants for whom eye data was obtained. (X-axis: left vs. right, $F(1,12) = 0.287$, $p=0.602$, up vs. down, $F(1,12) = 0.094$, $p=0.764$, interaction $F(1,12)=0.042$, $p=0.841$). (Y-axis: left vs. right, $F(1,12)=0.287$, $p=0.602$, up vs. down, $F(1,12)=0.094$, $p=0.764$, interaction, $F(1,12)=0.042$, $p=0.841$). There were no statistical differences in the number of data points collected for any of the conditions. As a result, the absence of systematic deviations from fixation precluded the need to exclude imaging runs.

3.3.2 Correlational Classifier using multiple features

I found above chance classification for all four attentional conditions in occipital ROIs, which included primary visual cortex (Table 3-1, Figs 3-2 to 3-5). Bilateral parietal, and bilateral LOC ROIs performed well, with LOC demonstrating above chance classification in all conditions except diagonal shifts of attention.

Table 3-1. Statistical results for correlational classifier.

Reports of means, standard deviations, and the results of t-tests comparing classification using the correlational classifier, against chance (25% for 4 way, 50% for dual choice conditions), for all four attentional conditions, in each ROI.

(Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom. Entries in red are statistically significant results, p< 0.05).

Attention Condition	Acc	SD	DoF	t- value	p-value
4 way					
Bilateral Occipital	32.69	5.85	12	4.74	0.00
Bilateral Parietal	28.69	5.46	12	2.44	0.02
Bilateral SMA	23.33	6.97	4	-0.53	0.69
Bilateral LOC	31.41	5.92	12	3.90	0.00
Bilateral FEF	26.49	7.09	6	0.56	0.30
Left vs. Right					
Bilateral Occipital	55.61	5.40	12	3.74	0.00
Bilateral Parietal	55.29	6.17	12	3.09	0.00
Bilateral SMA	50.83	6.69	4	0.28	0.40
Bilateral LOC	57.85	5.78	12	4.90	0.00
Bilateral FEF	51.79	8.11	6	0.58	0.29
Up vs. Down					
Bilateral Occipital	53.53	5.85	12	2.17	0.03

Bilateral Parietal	48.88	8.49	12	-0.48	0.68
Bilateral SMA	46.25	6.97	4	-1.20	0.85
Bilateral LOC	54.01	6.61	12	2.18	0.02
Bilateral FEF	50.30	3.49	6	0.23	0.41
Diagonals					
Bilateral Occipital	5 7.21	7.82	12	3.32	0.00
Bilateral Parietal	47.12	5.07	12	-2.05	0.97
Bilateral SMA	47.50	4.97	4	-1.12	0.84
Bilateral LOC	52.08	5.71	12	1.32	0.11
Bilateral FEF	46.73	5.88	6	-1.47	0.90

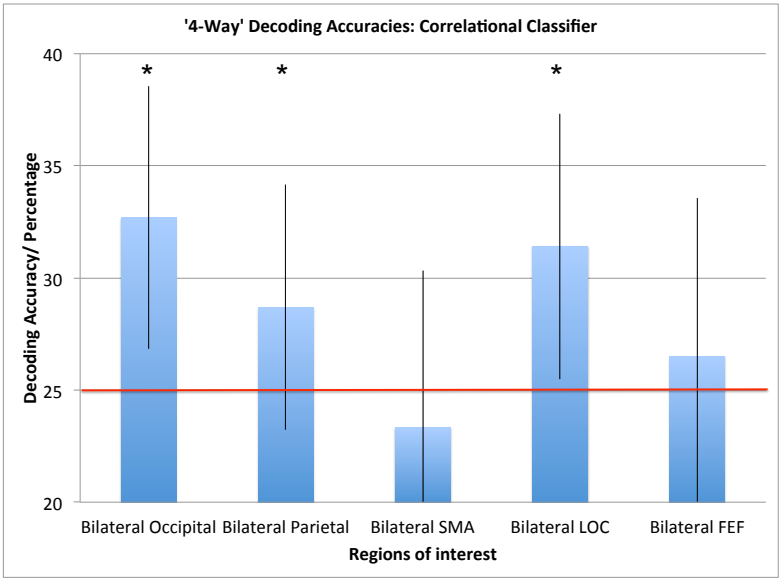

Fig 3-2. Four-way classification across ROIs with correlational classifier, averaged across participants. The stars over the bars denote statistically significant classification ($p<0.025$, one tailed t-test). The red line indicates 'chance' level classification at 25%. The error bars indicate ±1 SD.

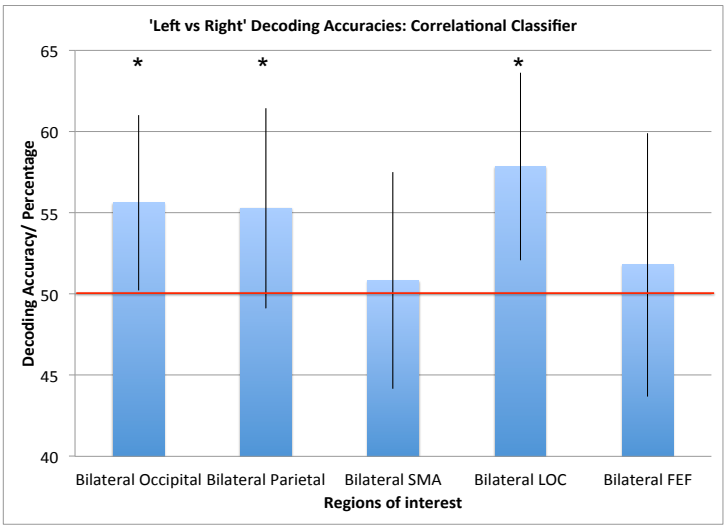

Fig 3-3. Left versus right classification across ROIs with correlational classifier, averaged across participants. The stars over the bars denote statistically significant classification ($p<0.025$, one tailed t-test). The error bars indicate ±1 SD. The red line indicates 'chance' level classification at 50%.

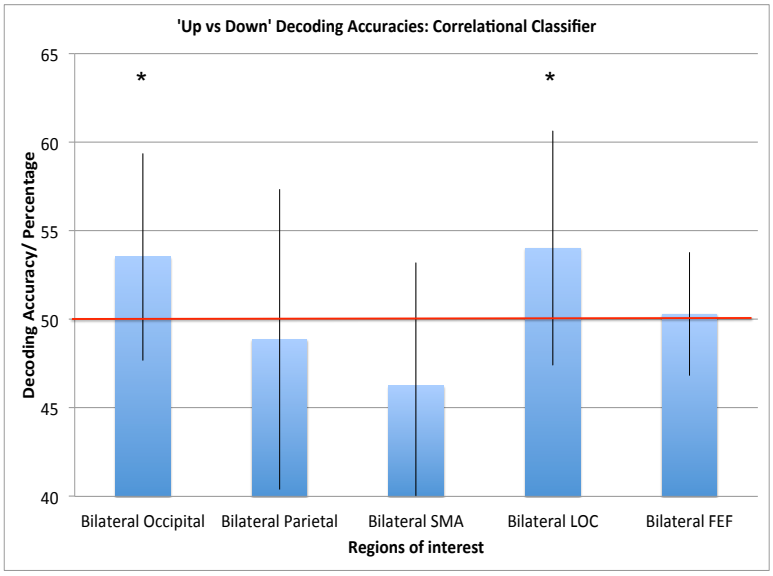

Fig 3-4. Up versus down classification across ROIs, with correlational classifier averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

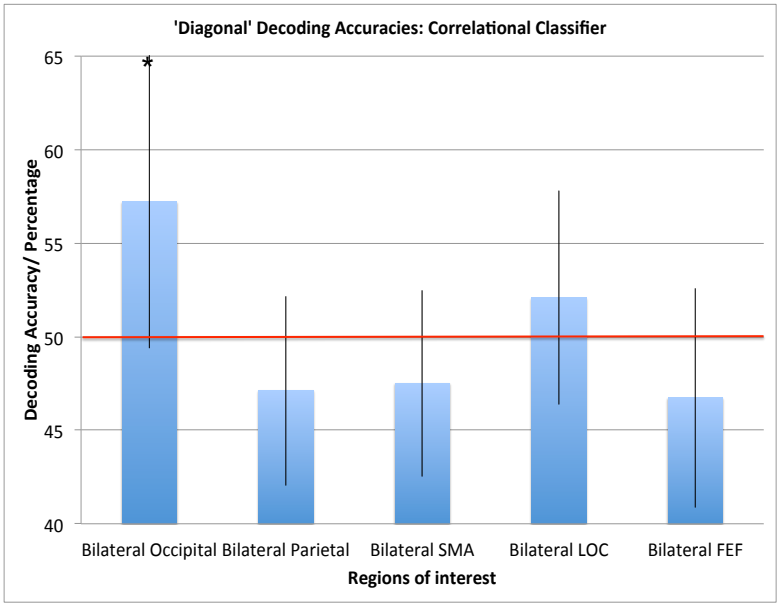

Fig 3-5. Diagonals classification across ROIs, with the correlational classifier averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

3.3.3 Linear Discriminant Analysis

Above chance classification only occurred in the 4-way and Left versus Right conditions, with Bilateral Occipital and LOC showing the highest classification accuracies (Table 3-2, Figs. 3-6 to 3-9).

Table 3-2. Statistical results for LDA classifier.

Means, SD, and the results of t-tests comparing classification with the LDA classifier, against chance (25% for 4 way, 50% for dual choice conditions), for all four attentional conditions, in each ROI. (Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom. Entries in red are statistically significant results, p<0.05).

Attention condition	Acc	SD	DoF	t- value	p-value
4 way					
Bilateral Occipital	30.61	7.29	12	2.78	0.01
Bilateral Parietal	24.20	5.68	12	-0.51	0.69
Bilateral SMA	21.67	4.80	4	-1.55	0.90
Bilateral LOC	29.65	8.55	12	1.96	0.04
Bilateral FEF	25.89	5.49	6	0.43	0.34
Left vs. Right					
Bilateral Occipital	53.85	5.23	12	2.65	0.01
Bilateral Parietal	53.37	5.48	12	2.21	0.02
Bilateral SMA	49.58	6.32	4	-0.15	0.56
Bilateral LOC	57.37	8.57	12	3.10	0.00
Bilateral FEF	49.70	4.57	6	-0.17	0.57
Up vs. Down					

Bilateral Occipital	52.40	7.55	12	1.15	0.14
Bilateral Parietal	48.88	8.62	12	-0.47	0.68
Bilateral SMA	52.08	4.42	4	1.05	0.18
Bilateral LOC	50.64	6.04	12	0.38	0.35
Bilateral FEF	50.60	7.19	6	0.22	0.42
Diagonals					
Bilateral Occipital	50.32	7.50	12	0.15	0.44
Bilateral Parietal	46.80	5.55	12	-2.08	0.97
Bilateral SMA	43.75	6.25	4	-2.24	0.96
Bilateral LOC	50.64	6.93	12	0.33	0.37
Bilateral FEF	51.19	4.79	6	0.66	0.27

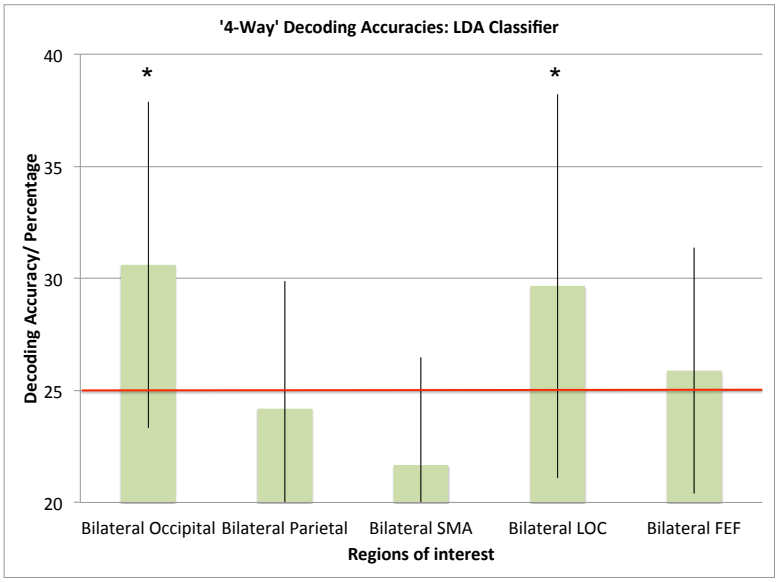

Fig 3-6. Four-way classification with LDA across ROIs, averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

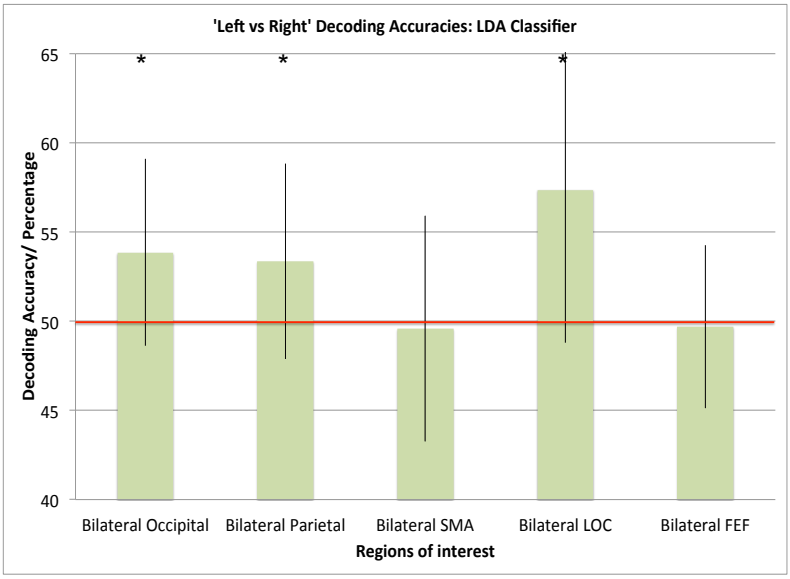

Fig 3-7. Left versus right classification with LDA across ROIs, averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

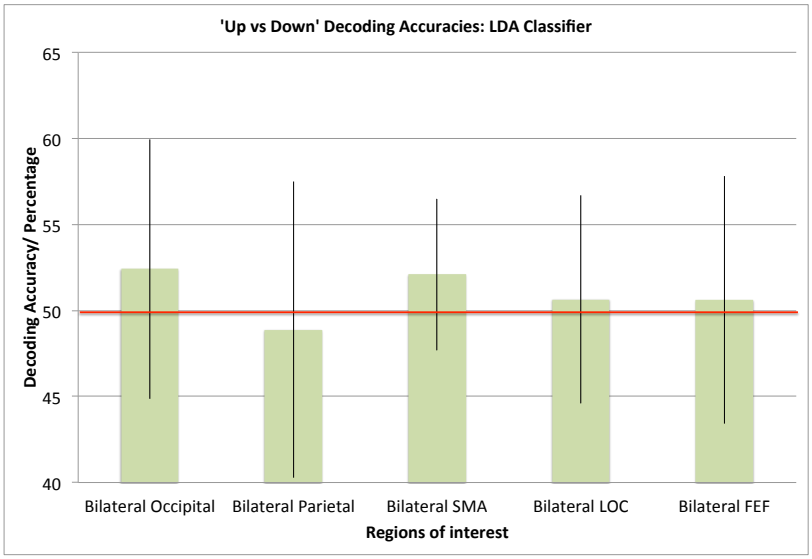

Fig 3-8. Up versus down classification with LDA across ROIs, averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

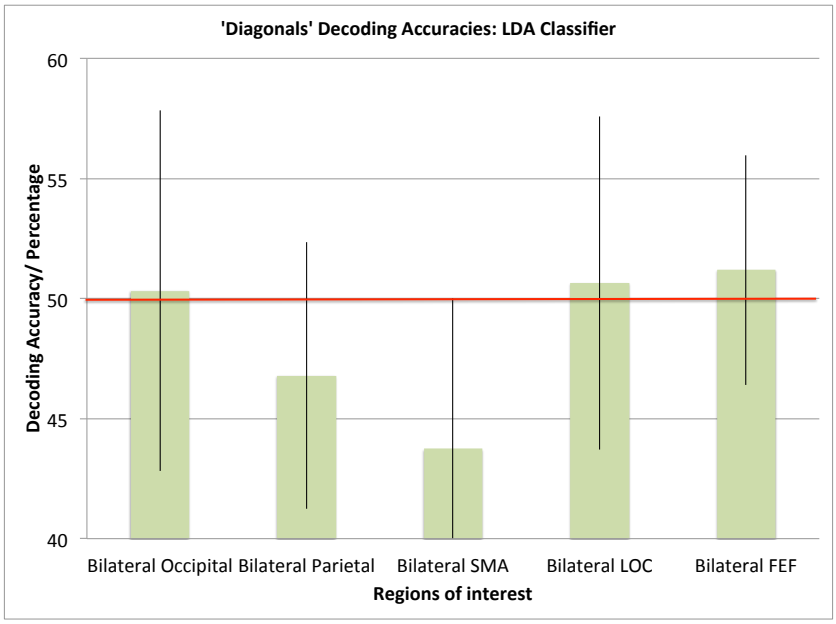

Fig 3-9. Diagonals classification with LDA across ROIs, averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

3.3.4 Linear Discriminant Analysis with Principal Component Analysis

Bilateral occipital ROIs performed well in all 4 conditions, with bilateral LOC performing next best (3 out of 4 conditions) (Table 3-3, Figs 3-10 to 3-13).

Table 3-3: Statistical results for LDA classifier with PCA.

Means, standard deviations, and the results of t-tests comparing classification with the LDA classifier using PCA, against chance (25% for 4 way, 50% for dual choice conditions), for all four attentional conditions, in each ROI. (Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom. Entries in red are statistically significant results, p<0.05).

Attention condition	Acc	SD	DoF	t-value	p-value
4 way					
Bilateral Occipital	29.97	5.74	12.00	3.12	0.00
Bilateral Parietal	26.92	4.54	12.00	1.53	0.08
Bilateral SMA	23.33	3.73	4.00	-1.00	0.81
Bilateral LOC	27.72	4.11	12.00	2.39	0.02
Bilateral FEF	24.70	1.44	6.00	-0.55	0.70
Left vs. Right					
Bilateral Occipital	55.13	5.35	12.00	3.45	0.00
Bilateral Parietal	52.24	4.46	12.00	1.82	0.05
Bilateral	47.92	2.95	4.00	-1.58	0.91

SMA					
Bilateral LOC	55.29	6.78	12.00	2.81	0.01
Bilateral FEF	51.49	2.88	6.00	1.37	0.11
Up vs. Down					
Bilateral Occipital	54.33	6.04	12.00	2.58	0.01
Bilateral Parietal	50.48	4.00	12.00	0.43	0.34
Bilateral SMA	51.67	3.42	4.00	1.09	0.17
Bilateral LOC	52.56	5.96	12.00	1.55	0.07
Bilateral FEF	52.98	2.03	6.00	3.87	0.00
Diagonals					
Bilateral Occipital	52.56	3.52	12.00	2.62	0.01
Bilateral Parietal	49.36	5.27	12.00	-0.44	0.67
Bilateral SMA	50.83	3.78	4.00	0.49	0.32
Bilateral LOC	51.60	5.32	12.00	1.09	0.15
Bilateral FEF	47.32	5.20	6.00	-1.36	0.89

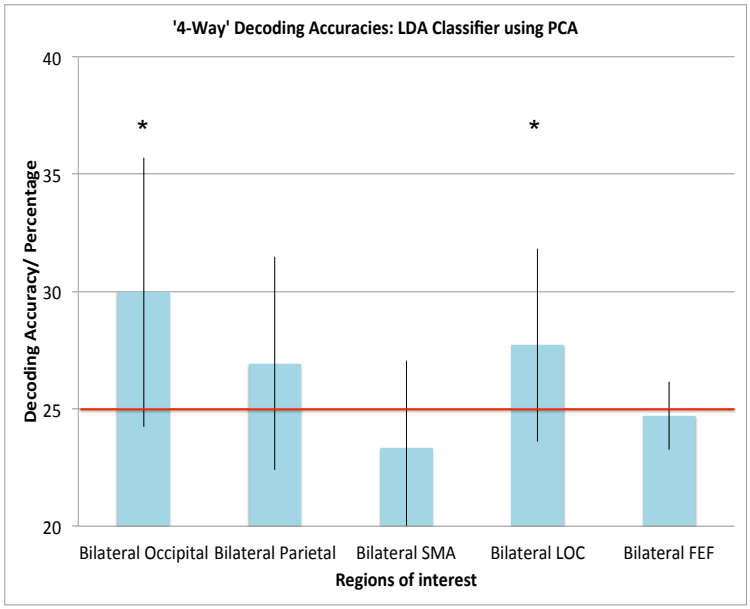

Fig 3-10. Four-way classification across ROIs for LDA with PCA, averaged across participants. The stars over the bars denote statistically significant classification ($p<0.025$, one tailed t-test). The red line indicates 'chance' level classification at 25%. The error bars indicate ±1 SD.

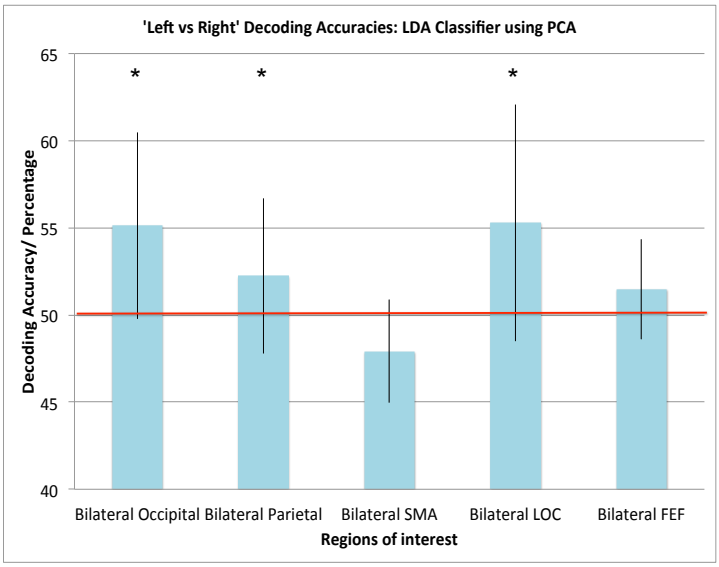

Fig 3-11. Left versus right classification across ROIs, for LDA with PCA, averaged across participants. The stars over the bars denote statistically significant classification ($p<0.025$, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

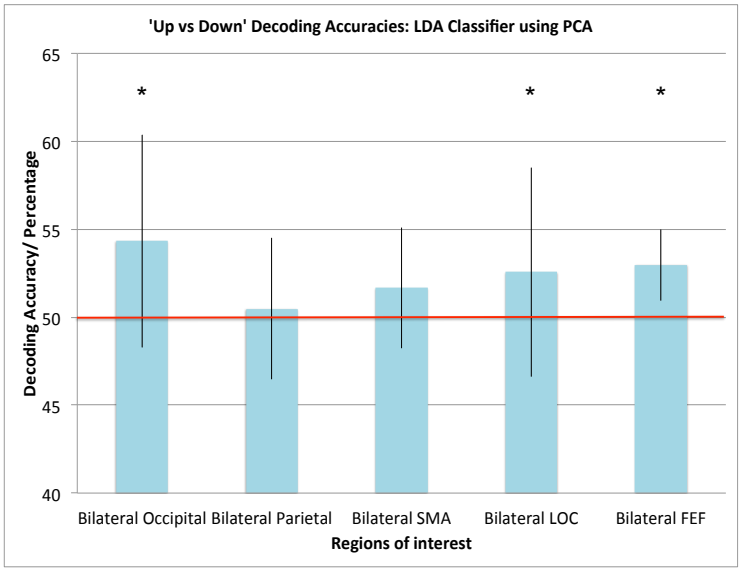

Fig 3-12. Up versus Down classification across ROIs, for LDA with PCA, averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

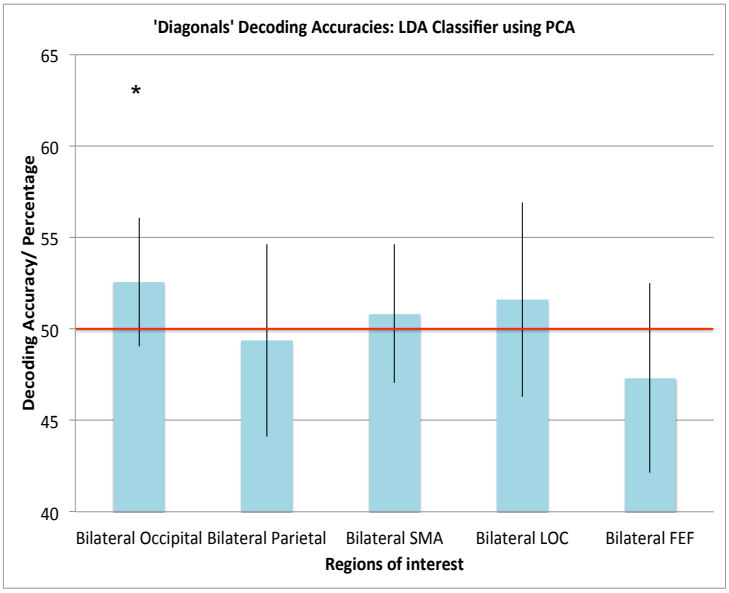

Fig 3-13. Diagonals classification across ROIs, for LDA with PCA, averaged across participants. The stars over the bars denote statistically significant classification (p<0.025, one tailed t-test). The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

3.3.5 Predicting shifts of attention using univariate approaches

The rationale for the use of univariate classification approaches has been outlined above. The results were less successful as compared with the multivariate classification accuracies. There was no significant classification from any of the ROIs for any of the attentional conditions using this approach.

Table 3-4: Statistical results for univariate classification.

Means, standard deviations, and the results of t-tests comparing classification with univariate classification (average of t-values), against chance (25% for 4 way, 50% for dual choice conditions), for all four attentional conditions, in each ROI. Nil significant results. (Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom).

Attention condition	Acc	SD	DoF	t-value	p-value
4 way					
Bilateral Occipital	24.04	4.30	12	-0.81	0.78
Bilateral Parietal	23.24	3.88	12	-1.64	0.94
Bilateral SMA	22.50	9.13	4	-0.61	0.71
Bilateral LOC	24.68	5.02	12	-0.23	0.59
Bilateral FEF	23.81	3.78	6.	-0.83	0.78
Left vs. Right					
Bilateral Occipital	51.92	7.63	12	0.91	0.19
Bilateral Parietal	52.24	7.58	12	1.07	0.15
Bilateral SMA	43.33	5.97	4	-2.50	0.97
Bilateral LOC	50.00	6.01	12	0.00	0.50
Bilateral FEF	51.79	5.30	6	0.89	0.20
Up vs. Down					
Bilateral	50.48	8.30	12	0.21	0.42

Occipital					
Bilateral Parietal	50.16	6.22	12	0.09	0.46
Bilateral SMA	50.42	6.49	4	0.14	0.45
Bilateral LOC	51.76	5.57	12	1.14	0.14
Bilateral FEF	52.38	6.53	6	0.97	0.19
Diagonals					
Bilateral Occipital	48.56	4.76	12	-1.09	0.85
Bilateral Parietal	46.15	6.58	12	-2.11	0.97
Bilateral SMA	46.25	2.28	4	-3.67	0.99
Bilateral LOC	48.56	6.33	12	-0.82	0.79
Bilateral FEF	52.08	5.77	6	0.96	0.19

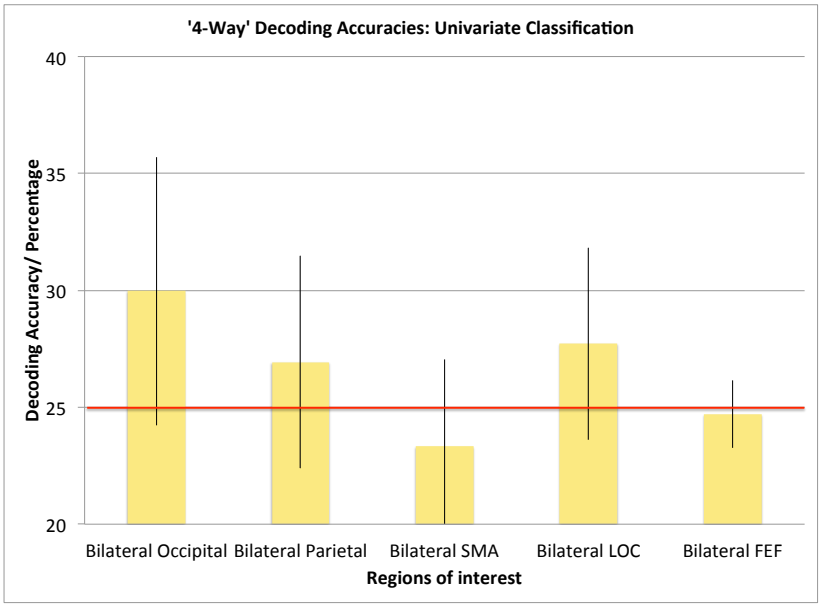

Fig 3-14. Four-way classification across ROIs, averaged across participants, for Univariate Classification. The red line indicates 'chance' level classification at 25%. The error bars indicate ±1 SD.

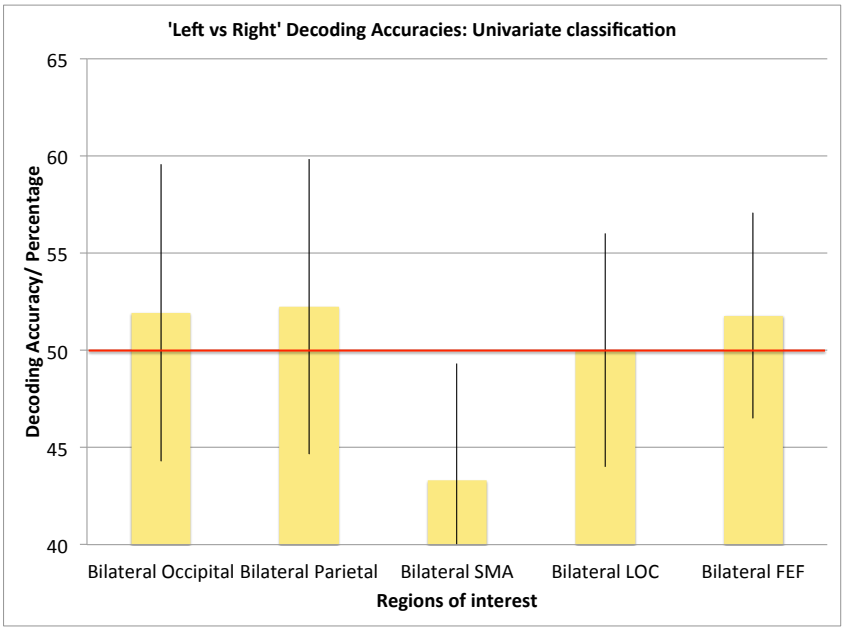

Fig 3-15. Left versus right classification across ROIs, averaged across participants for Univariate Classification. The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

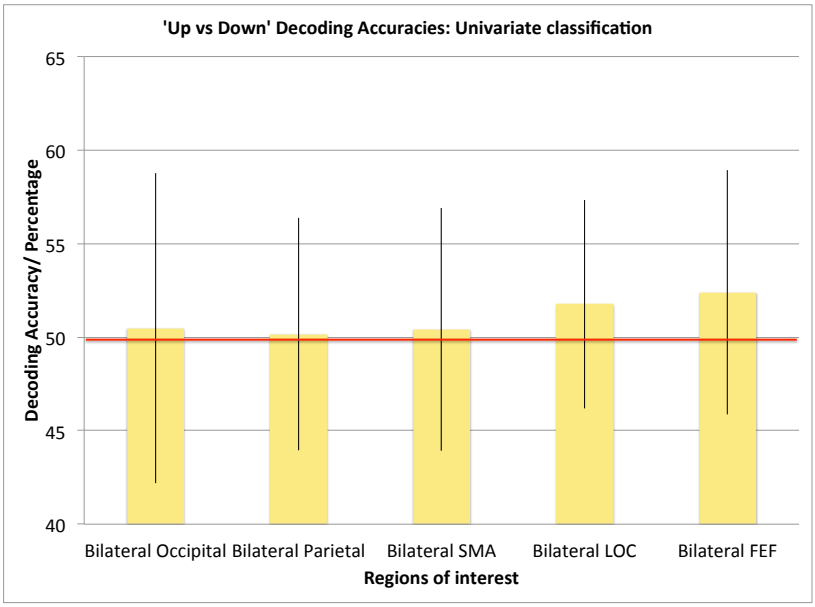

Fig 3-16. Up versus down classification across ROIs, averaged across participants for Univariate Classification. The red line indicates 'chance' level classification at 50%. The error bars indicate ±1 SD.

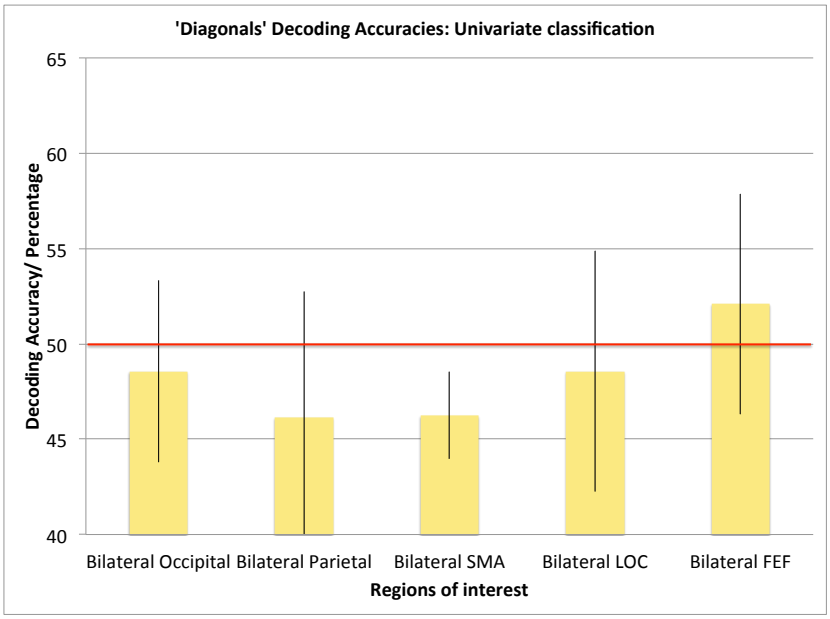

Fig 3-17. Diagonal classification across ROIs, averaged across participants for Univariate Classification. The red line indicates 'chance' level classification at **50%. The error bars indicate ±1 SD.**

3.3.6 Correlational classifier with the ‘single best voxel’

The use of the correlational classifier with the ‘single best voxel’ was attempted but yielded no significant results in terms of classification for any ROI, for any attentional condition.

3.4 Discussion

I explored whether the location of covertly deployed spatial attention could be decoded from activity recorded non-invasively by fMRI from higher order visual areas. I used two approaches (multivariate and univariate classification) on a spatial attention task performed during fMRI. My hypothesis was that spatial attention deployment could be successfully decoded from activation in higher order visual cortex, with accuracy levels comparable to primary visual cortex. I predicted the use of multivariate classification techniques would yield highly accurate classification from these regions. An open question was how well univariate classification would perform in comparison. I found highest rates of classification for all attention conditions with multivariate classification specifically in bilateral occipital regions (including primary visual cortex) and bilateral LOC. Univariate classification was not successful at decoding the location of spatial attention. Overall, consistent classification was less reliable in frontal and parietal regions.

This investigation was based on two existing observations in the literature. Firstly, the presence of topographic maps of visual space throughout visual cortex, including both primary visual cortex and higher order regions such as the parietal lobe and the LOC. Secondly, the co-existence of 'attentionotopy', that is the preferential enhancement of these retinotopic maps in relation to deployment of covert spatial attention, particularly within higher order visual areas in frontal, parietal, and lateral temporo-occipital regions (41,374,379).

Primary visual areas respond robustly to high contrast visual stimuli such as checkerboard gratings, driving activation of appropriately specific retinotopic co-ordinates of visual space. The effect of attention on these regions is well-documented (119,394–397). In this study the stimuli were of sufficient contrast (i.e. black and white spirals) to elicit a salient response in primary visual cortex. Therefore classification from primary visual ROIs served as a comparator for classification accuracies in higher visual areas for the effect of attention on the underlying retinotopy in these regions.

Studies demonstrating retinotopy and 'attentionotopy' in high order visual areas such as in the parietal lobe along IPS (398,399) (saccade driven) FEF and inferotemporal cortex (41) have successfully used univariate analysis techniques with arguably more salient or biologically relevant stimuli compared to our own study which used small white spirals (2 degrees visual angle) on a black background. Silver's landmark experiment utilised contrast gratings in a 'piechart' of spatial locations (398), while Saygin and Sereno used rotating wedges containing point-light biological motion figures (41). It is likely that the attention 'signal' in higher order visual cortex incorporates the biological importance of the stimulus being attended to, as well its visual salience, thus incorporating top-down as well as bottom-up influences (362). It may be for this reason that I did not find consistent classification in some of the higher order regions such as parietal cortex and FEF.

I included SMA as a ROI as I speculated that it might have a top-down role in sustained visuospatial attentional allocation. The SMA has a putative role in cognitive control in the presence of conflict resolution, and in relation to internally generated or self-initiated action (400). On this basis I attempted to classify the endogenous shifts of attention set out in our experiment task. No above chance classification took place with any of the classifiers. A possible reason for this is that the functional involvement of SMA is more specifically related to its role in the volitional component of self-directed action rather than the trigger for attentional allocation.

An interesting finding was the consistently high classification accuracy for attentional conditions in the LOC. It has previously been suggested that the LOC may serve as an intermediate hub in the top-down attentional network, with the presence of both object-sensitive and retinotopic information here (381). This will be discussed in more detail below.

The specific deployments of attention that I examined were 'left versus right', 'up versus down', diagonal and four-quadrant attention. I used comparatively simple but robust multivariate classification algorithms (correlational classifier, LDA, and LDA using PCA), with t-maps for single voxel estimates (269). Across the three multivariate classifiers used, significantly above chance classification in

distinguishing ‘left versus right’, and for the four quadrants, was observed in primary visual areas (bilateral occipital), and to a lesser extent in higher visual areas. Bilateral LOC was as successful for these two conditions, across all three classifiers. Classification from within the bilateral parietal ROI was above chance with the correlational classifier, and the LDA classifier using PCA, for both conditions. None of the other regions examined provided above chance classification. Examining ‘diagonals’ and ‘up versus down’ conditions provided above chance classification only with the bilateral occipital ROI, and then only with the correlational classifier, and LDA using PCA. ‘Up versus down’ with the correlational classifier, and LDA with PCA could successfully decode spatial attention deployment from the Bilateral LOC ROI. Bilateral FEF only provided accurate classification for ‘up versus down’ using the LDA with PCA approach.

No accurate classification was achieved with the LDA classifier (i.e. without prior PCA). None of the other regions studied provided accurate classification for the ’diagonals’ condition.

My results indicate that using higher cortical areas to establish the current focus of spatial attention in either upper or lower visual fields is challenging. With regards to spatial attention, Kraft et al. have demonstrated that spatial orienting has a lower visual field preference (401). Moreover, shifts of attentional focus are more precisely delineated when performed in the upper visual field, whilst stationary attention has a lower visual field advantage (402–404). Discerning ‘up versus down’ and ‘diagonal’ deployments of attention were more challenging to decode in the ROIs. This may be related to the asymmetric cortical representation of retinotopy with regards to the upper and lower visual fields, resulting in a differential distribution of the effects of attention (401,405). Despite the documented presence of retinotopy in the parietal cortex, LOC, and FEF (see Introduction), it has been suggested that areas of the visual field are asymmetrically represented in the higher order visual areas, and that there is a visual field bias which varies from brain region to brain region. While there is physiological evidence for lower field preferences at the level of the retina (as much as 60% of retinal area) (406) and the lateral geniculate nucleus (407), the correlates of cortical asymmetries are incompletely understood. Monkey studies have demonstrated lower visual field advantages in higher visual areas (V6a, middle

temporal area (MT)) (408,409). More recently, Sayres and Grill-Spector et al. (98,410) have demonstrated lower visual field bias, in the presence of complete retinotopic representation in the LOC. Schwarzlose et al. reported an upper visual field bias in the PPA and a lower visual field bias in FFA (98).

Using the univariate approach of averaged t-values, no significant classification accuracies were achieved in any of studied regions. Additionally, the classification obtained using a single 'best' voxel was also not successful. The reason for this finding with the univariate approaches may be linked to several factors. The comparative simplicity of the stimuli used i.e. spirals, as compared to high contrast gratings, may have resulted in lower magnitudes of BOLD activation through the visual hierarchy. As such, fine grained activations related to the stimuli and the deployment of spatial attention would be more likely to be identified by multivariate approaches. Primary retinotopic areas are highly sensitive to the physical properties of the stimulus, such as contrast, edges, spatial frequency, and also have comparatively lower levels of BOLD activation in response to visual stimuli as compared to higher visual areas. This may explain the lack of success in classification using univariate approaches in these regions. As established earlier, higher order visual areas respond to more complex visual features of a visual stimulus such as category-specific information and semantic content. The activation produced in these regions in response to simple, visual spirals may have not been of sufficient magnitude therefore, to be identified by univariate approaches.

The success with the pattern correlational classifier and LDA with PCA may be related to a number of factors. They are computationally inexpensive, simple and rapid to train. In particular, the correlational classifier is appealing with regards to its straightforward interpretation i.e. it provides the most correlated class-averaged training pattern. Simple linear classifiers perform well in the presence of smaller data sets (e.g. hundreds of voxels rather than thousands of voxels); this would also explain the strength of classification with the LDA classifier when using eigenvariates produced by a PCA. The latter enables a dimensionality reduction, potentially increasing the signal to noise ratio. Further the true distributions of the Bayes-optimal (i.e. ignoring estimation error) decision boundaries may have been approximately linear. It is likely that the task, presenting visual-spatial information utilising

retinotopic representations, generates linearly classifiable information that is optimally suited to the use of these classifiers. Correlational classifiers classify patterns in keeping with their Pearson correlation coefficient with a category template pattern. They utilise a nearest neighbour approach, with within-class averaging and a correlation distance function. The test pattern is classified according to the category pattern it is most similar to. On the other hand, the LDA classifier utilises a Gaussian approach, with a Mahalanbois distance function. It determines the discriminant dimension in response-pattern space, maximising the ratio of between-class over within-class variance of the data. As such I provide evidence for the use of simple linear classifiers for multivariate classification of brain activation in relation to the deployment of spatial attention.

I used classification accuracy as the means of assessing and reporting BCI performance, and compared it with classification at chance. This was motivated by classification accuracy directly corresponding to the probability of performing a correct classification. Other comparable fMRI BCI studies have reported BCI performance in terms of classification accuracy (43,44). A necessary caveat for the use of classification accuracy is that classes for classification must be balanced, as was the case in this study. Alternative metrics include sensitivity and specificity, which measure the proportion of correctly identified positive results (true positives) and the proportion of correctly identified negative results (true negatives). Sensitivity is alternatively referred to as the ‘true positive rate’ or ‘recall’, and bears a relationship to classification accuracy. Other measures which have been used include F-measures which examines precision and recall, and correlation co-efficients which are used for the validation of classification (411). I used t-tests to determine whether classification accuracy was significantly different from chance-level classification, an approach that can be considered a standard way to assess classifier performance (e.g.(412–414)).

In this study I explored classification algorithms that could be simply incorporated into an attention-driven BCI. Although I achieved above chance classification accuracies in the attention conditions, the magnitude of ‘decoding’ accuracy was comparatively modest in the context of an application for a communication interface. Current recommendations suggest target BCI accuracies in the region of 70-80%

(78), (this recommendation was applied broadly to all forms of BCI i.e. EEG). However the influence of cognitive task on classification accuracy as well as, individual differences in brain activation responses, and imaging modality, remains unknown. Nonetheless, I consider further methods for increasing information transfer rate, and the production of higher rates of classification accuracy in the next chapter.

I attempted to identify higher order brain regions that would provide optimal information for classifying the direction of visuospatial attention. It would appear from my work, that higher order visual regions provide the opportunity for further exploration, by investigating the use of a more detailed neural signal, incorporating information related to object category, as well as spatial information. From my results, LOC appeared to be the most promising cortical target in this regard. This was considered an important finding from a biological perspective, in that LOC is traditionally recognised as having a role in object- specific attention. Its role in attention in relation to spatial position is less clearly understood. There is a growing body of literature examining the presence of both object-based and spatial information in the LOC (48,112,383,384,415). A number of authors have demonstrated position related information in category-specific cortical regions such as LOC (410,416,417). Arguments for a spatiotopic 'grid', indicating the positioning of the same object at different locations in space have therefore been made, although evidence to the contrary has also been shown, suggesting instead that object sensitive information is position tolerant (48). Importantly, Carlson et al. (410,416,418) have suggested that spatiotopic information may not be accessible to active processing. It is possible that the presence of object and spatial information in the same cortical region may be encoded by neighbouring but discrete populations of neurones, which cannot be measured with the spatial resolution currently afforded by fMRI.

An important consideration in this study, given the examination of multiple brain regions, as well as a number of classification approaches, is the problem posed by multiple comparisons. Typically, a Bonferonni correction can be applied, which attempts to establish the family-wise error rate, that is the likelihood of falsely reporting a significant classification (i.e. a Type 1 error), when multiple brain regions are being tested simultaneously. For n comparisons, the significance is adjusted by

1/n. Although applying this correction to the classification results obtained would not effect all of the significant results presented in this study, Bonferroni corrections in the context of BCI bio-signal processing has been suggested to be too conservative (411). This relates to the fact that it assumes independence of comparisons, when BCI classification is frequently performed on co-dependent temporal and spatial data. The false discovery rate (FDR) has been proposed as an alternative, controlling the proportion of false positives, rather than the probability of a single false positive. However, the FDR increases the likelihood of Type 1 errors. A useful approach may instead be to adopt a hierarchical significance testing approach. This has been used in EEG analysis (419), and may be applied in principle to fMRI analysis. In this case, the data is put into 'family hierarchy', with sub-families (child hypothesis) i.e. a family of parietal ROIs. Child hypotheses are compared recursively if 'their parents' null hypothesis is rejected. This experiment (and experiment 2) were fundamentally aimed at exploring potential approaches to classifying spatial attention, and by necessity tested several classifiers and brain regions.

In light of my findings, and previous work (48,410,416,417) LOC may potentially serve as a preparation platform for the integration of information from both dorsal and ventral visual streams. The strength of the BOLD signal, measured in the LOC in relation to the deployment of spatial attention may be enhanced through the use of more complex stimuli such as real-world objects. Object-selective cortex which demonstrate retinotopy such as LOC might therefore be suitable targets for BCIs which use neural activity associated with real-world stimuli at differing spatial positions of the visual scene(15,42,43,45,323,420). Such a BCI could be used by paralysed patients who might be unable to speak, but wish to interact independently with the environment (421). The device would need to be intuitive to use, but enable interactions with real-world objects. A practical example demonstrating the ultimate evolution of this approach would be a bed bound patient interacting with a bedside visual interface powered by his decoded neural activity. The interface would be a partitioned screen, with multiple object streams, which might include food items, faces of carers, and body parts. As required, the patient would direct attention to one of the spatially separated streams, using his neural activity to 'select' a relevant stimulus i.e. an image of his back to indicate back pain, and image of glass of water to indicate thirst etc. A realtime fMRI-based BCI could be used as the first step in the

development of a communications platform, serving to identify and hone specific neural activity and brain regions, prior to conversion to an operational implantable BCI. This concept extends attention-driven BCIs beyond being driven by purely retinotopic, primary visual cortex based neural activity (10,44), to the use of higher order visual regions, with 'decoding' being applied to stimuli with real-world relevance. I explore this further in the next chapter.

3.5 Conclusion

The application of multivariate analyses in the presence of simple stimuli enabled the reliable identification of subtle patterns of spatial information and indicated where spatial attention was being deployed by a participant. This was achieved in a number of previously identified cortical regions involved in spatial attention. I additionally provide evidence for the presence of spatial information in the LOC. While this region is known to have both spatiotopic and object-sensitive information, this is, to my knowledge, the first study illustrating the involvement of the LOC in the deployment of spatial attention utilising MVPA.

3.6 Rationale for Experiments 1 and 2

The aim of my first experiment was to identify a usable fMRI paradigm in which the deployment of spatial attention could be used to develop a rt-fMRI BCI. Prerequisites for such a device include simplicity of implementation, ease of use, and real-world applicability, in a similar manner to the choice of multivariate classifier used in the data analysis. In my first experiment I identified a number of higher order visual areas as well as primary retinotopic cortex that classified the direction of spatial attention. This adds to recent literature documenting the use of classification techniques to determine the direction of spatial attention (10). In the next experiment I extend these findings by examining the deployment of spatial attention to stimuli with real-world salience. The use of such stimuli is of operational use in the context of a BCI in a clinical setting, for example in use by locked-in patients. Further, I hope

to take advantage of object-based information, in addition to spatial information in higher order visual cortex, for the purposes of ‘decoding’ the direction of attentional deployment.

ABSTRACT: Experiment 2

Predictive encoding of spatiotemporal patterns in higher order visual areas using real world stimuli- multivariate and univariate approaches

Introduction

Higher order visual areas such as LOC contain information about both objects and their location. Here I test the hypothesis that the deployment of visual attention to different object categories in specific regions of the visual field can be ‘decoded’ using univariate and multivariate classification applied to BOLD signals from higher order visual areas, as well as from early retinotopic visual areas.

Methods

Participants were scanned performing a task in which 4 stimuli were presented simultaneously in each of four visual quadrants. Stimuli consisted of 4 categories (faces, houses, body parts, and consumables). Participants were instructed to fixate centrally and voluntarily attend to one of the four quadrants. At the end of a block they disclosed which quadrant they had previously been attending.

Within functionally defined ROIs, representing retinotopic and higher visual areas responsive to the stimuli, univariate classification employing a 'winner take all' decision rule was used to blindly determine which quadrant had been attended. This prediction was compared against participants' behavioural responses to generate a measure of online prediction accuracy across blocks and for each ROI. I further compared the performance of this simple decision rule against multivariate classification accuracy.

The decision rule produced significant classification accuracies in all the ROIs except one (SPL) ranging from 25% to 50% (chance = 25%).Although not as successful, multivariate classification was most robust with linear discriminant analysis.

Classification accuracies ranged from 25% to 30% with chance level at 25% and reached statistical significance only in a minority of participants.

Conclusions

Classifying the direction of spatial attention using simple univariate approaches on a single block basis was achieved significantly better than chance from higher visual areas including LOC and FFA. The use of multivariate classifiers also demonstrated above chance accuracy in similar higher order visual areas but did not improve upon classification accuracy. Classifying spatial attention in LOC, including temporal information and stimuli with real-world saliency may pave the way towards operational applications such as 'attention-based BCIs' for locked-in patients.

Key words

Realtime fMRI (rt-fMRI), spatial attention, classification, decoding, m-sequences, brain-computer interface (BCI), higher order visual areas, LOC, FFA

4. EXPERIMENT 2

Predictive encoding of spatiotemporal patterns in higher order visual areas using real world stimuli- multivariate and univariate approaches

4.1 Introduction

Real-world human actions are often preceded by the direction of attention to intended actions. This often begins in an abstract manner, with the interaction of cognitive functions directed by the overarching involvement of attention. The study of visual attention over the past three decades using fMRI has revealed the manner in which cortical retinotopic mapping of a visual scene can be modulated by attention (105,422–424). Retinotopic maps in higher order visual areas e.g. FEF, parietal lobe (425), LOC (112) FFA (41,410) have been shown to interact with the deployment of spatial attention. These maps have eccentricity representations as well as orderly polar angles, as seen in early visual cortex (112,371). Support for a hierarchical top-down organisation of attention (119,426–428), has led to the proposal of 'attentionotopy', in which discrete topographically-specific command signals from higher order visual areas, pass in a co-ordinated fashion through a hierarchical cortical circuit of attention (116), acting on specific retinotopic neuronal populations.

Attention may be allocated to space as well as to specific features and objects. This may occur within a circumscribed area of space, the so-called 'spotlight' of attention (357) , or throughout the visual scene as in the case of attending to a specific colour, or the motion of objects (429,430). The presence of both object and spatial information within the same regions of higher order visual areas, such as LOC and

FFA, has been demonstrated using multivariate and univariate analysis (98,417). An on-going debate is whether this information is segregated, or represented independently within the same neuronal population (431). The interaction of object-based attention in space may enable extraction of higher rates of information from functional imaging data for the classification or 'decoding' of the direction of spatial attention. As discussed in the previous chapter, objective-selective cortex which demonstrate retinotopy may be suitable for 'attention-based' brain-computer interfaces (BCI) (10,15,274,311), enabling neural activity linked to real-world stimuli at specific spatial positions in the visual scene to be used as a surrogate for explicit communication. This would allow patients to interact with their environment using their 'decoded' neural activity, despite being unable to speak, move their eyes, or move their limbs as is the case with 'locked-in' patients (2,4,15,421,432).

Here I explore the classification of the BOLD signal in relation to voluntary (i.e., not cue directed) deployment of spatial attention to four spatially separated streams of biologically meaningful stimuli. The aim of this chapter is to demonstrate that classification can be reliably performed on a trial-by-trial basis. In addition to using conjunctions of spatial and object-based attention, I introduce specific alterations of the temporal presentation of the stimuli, incorporating a further layer of information into the resulting fMRI BOLD signal. I aimed for maximum orthogonality between attended quadrants by using separate m-sequences (*maximum shift L-level register sequences*) for the timing of object presentation in each quadrant. M-sequences are pseudorandom sequences of integers that have recently been introduced in fMRI experimental design, to optimise stimulus presentation (249). The latter is effectively an extension of the concept behind event related fMRI i.e. iimproved resolution of haemodynamic responses associated with transient neural activity linked to stimulus presentation or task performance.

I use an fMRI attention paradigm, in which the participants attended to one of four quadrants showing streams of objects. A particular quadrant is covertly selected on a block-by-block basis. I therefore test the hypothesis that the deployment of visuospatial attention can be decoded on a single block basis using the BOLD signals produced in higher order visual areas, aided by using differences in stimulus category, timing and location.

4.2 Methods

4.2.1 Participants

Nine neurologically healthy adult volunteers (24–32 years of age; mean age = 28; 4 females) with normal or corrected-to-normal visual acuity were recruited from the general population to participate in the experiment. Each participant was provided with written informed consent approved by the local ethics committee (UCL Ethics Committee code: 09/H0716/14) and passed MRI safety medical screening approved by the Wellcome Trust Centre for Neuroimaging (where MRI scanning was performed).

4.2.2 Equipment set-up

Stimuli were back projected (NEC Lt Projector) onto a screen mounted to the top of the magnet bore behind the participant's head. All visual stimuli described below were generated and displayed via MATLAB scripts created with Cogent extensions. Participants viewed the screen (via a mirror mounted to the head coil) at an optical distance of 52cm. The participants responded via a pair of custom-built, MR-compatible, fiber-optic, push button response boxes.

4.2.3 Eye-tracking

Eye position and pupil diameter were sampled at 60 Hz using long-range infrared videooculography (ASL 504LRO Eye Tracking System, MA). Throughout the experiment eye movements were monitored on-line via a video screen, for all participants.

4.2.4 Stimuli and presentation procedure

The stimuli were drawn from 4 categories- faces, houses, body parts, and consumables. Each category contained 16 unique exemplars that would also be specific to each quadrant. In each task block, images (diameter of stimuli: 2 degrees, eccentricity: 6 degrees) from each object category were presented in a pseudo-random manner, consisting of 16 images per mini-block (400ms per image, 100ms gap); each task block was composed of 2 mini-blocks, such that all 16 exemplars within an object category were shown during a mini-block. The order of the category mini-blocks was counter balanced between and across sessions. All images were rendered with identical greyscale values, and mean luminance using a custom designed MATLAB script.

4.2.5 Stimuli for retinotopy

To map retinotopic visual areas, flashing checkerboards (8 Hz frequency) were simultaneously presented in all four quadrants for 30s, with rest intervals of 30s. This was repeated 4 times within a retinotopic mapping scanning run.

A second retinotopic localiser was used to functionally define the anatomical boundaries of primary visual cortex. This involved the presentation of a stimuli simultaneously consisting of a rotating checkerboard wedge, and an expanding/ contracting checkerboard ring. Presentation parameters were identical to those used for the first retinotopic localiser run.

4.2.6 Experimental Design

The experiment was divided into 5 sessions. Each session consisted of 8 blocks. Each block comprised of a cue interval (3s), the task (32s), and a rest period (12.52s). Participants were instructed to maintain central eye fixation at all times. This was facilitated by the presence of a dot at the centre of the screen. During the cue interval this was orange (‘get ready’), white during the task, and red during the rest period.

The structure of the experimental task varied between sessions. During the first session, streams of stimuli were presented in one quadrant of the screen for the duration of one block, such that each quadrant was attended to twice over the course of a session. This functional localiser session helped identify object-selective and spatial-specific regions of primary and higher order visual cortex. (Fig. 4-1, top panel.)

During the next 4 sessions, stimuli were presented simultaneously in all four quadrants. More specifically during the 2nd of the 4 sessions, attention to a particular quadrant was indicated using a directional cue, presented during the cue interval. Once again each quadrant was cued for attention twice over the course of the session. (Fig. 4-1, middle panel.)

During 3rd to 5th sessions, participants were instructed to covertly attend a quadrant of their choice for the duration of a whole block. Further they were advised to use a strategy that enabled them to attend all quadrants twice during a session. They disclosed the attended quadrant by means of a button press during the rest period at the end of each of block. (Fig. 4-1, bottom panel.)

4.2.7 Retinotopy

A separate retinotopy scanning session was performed over 45 minutes. This involved the use of two localiser sessions. A retinotopic localiser targeting the specific eccentricities of images used in the main experiment was used, with the stimuli matching the size and shape of the stimuli used in the main experiment, so as to produce matched retinotopic activations in primary visual cortex.

4.2.8 N-back task

In order to ensure attention was maintained in all sessions there was a button press required if two successive exemplars were identical (one-back task). This occurred between one to three times per mini-block.

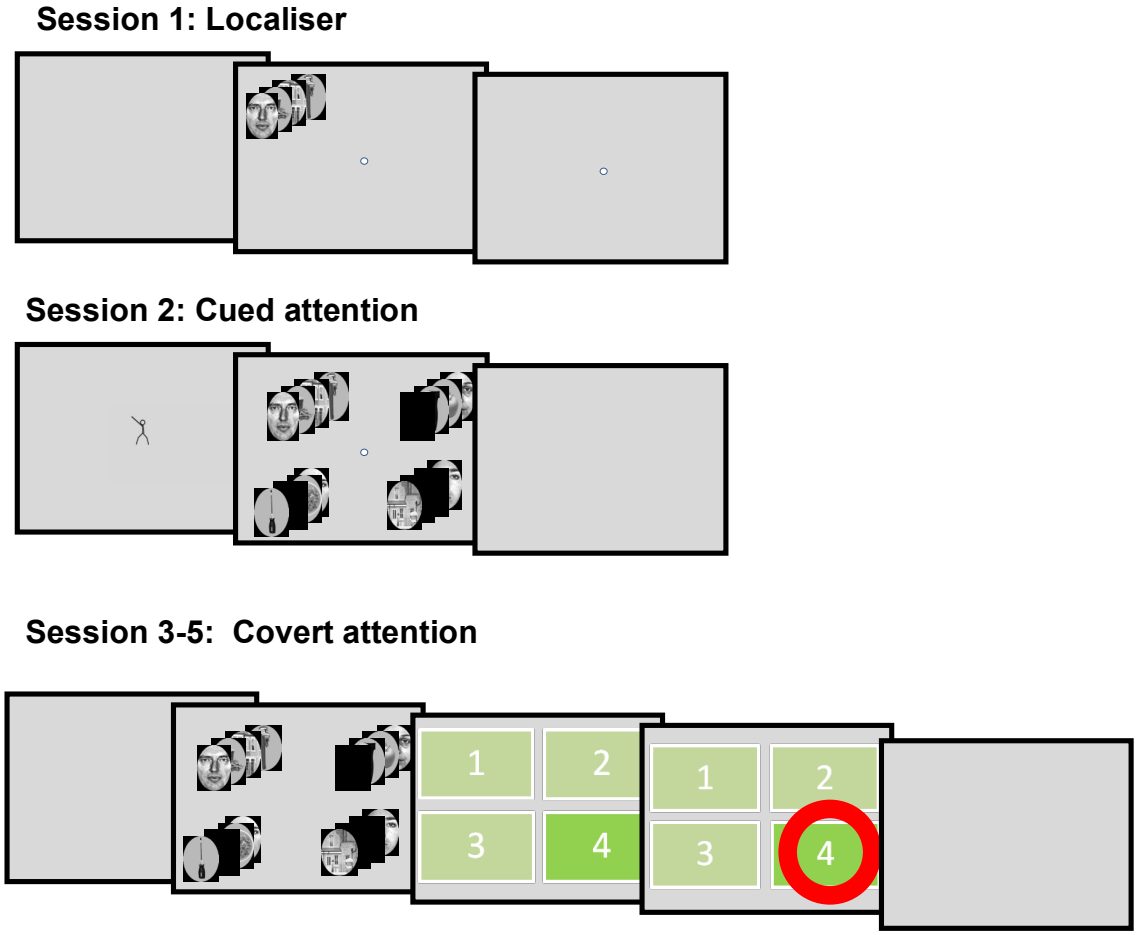

Fig 4-1. Experiment 2 design schematic.

Session 1 (top panel): Object/location functional localiser. Stimuli were only presented in one quadrant per block of attention; this was repeated twice. The purpose of this session was to clearly delineate functional activations associated with each of attentional conditions.

Session 2 (middle panel): Cued attention session. Same-category stimuli were simultaneously presented to all four quadrants. Stimulus frequency and category were unique to each quadrant in a similar fashion to session one. Attention was cued through the use of stick man during the cue interval indicating which quadrant was to be attended over the course of a block.

Sessions 3-5 (bottom panel): Voluntary attentional deployment. Participants were required to covertly direct their attention to a self-chosen quadrant for the length of one block. During the inter-block rest interval they revealed their chosen quadrant through a button box press linked to numbered visual interface- this was recorded to enable subsequent analysis. Each quadrant was attended to twice during the course of a session.

4.2.9 M-sequences, and modelling its effects on the haemodyamic response in a 4 quadrant spatial attention experiment

The timing of stimulus presentation was achieved through the use of four quadrant specific m-sequences. As described in chapter 2.4.1, m-sequences specify an optimal order of events that allows for the statistical disambiguation of different events presented close together in time, which would otherwise be difficult given the slow nature of the haemodynamic response (see chapter 2.2.2.1). This is achieved in two ways- m-sequences are nearly orthogonal to cyclically time-shifted versions of themselves, and they maximize the number of presentations for all event types.

The m-sequences provided 32 stimulus presentation slots per quadrant per block. Each of the 4 m-sequences were created to be maximally uncorrelated with each other. In order to test that the m-sequences for the four quadrants were optimally orthogonal to each other, a computational model of simulated brain activations were examined by taking each m-sequence and convolving it together with modelled ‘noise’ and the HRF.

After generating four m-sequences that were uncorrelated, and convolving them with a canonical HRF, I sampled the frequencies down to the typical TR (i.e. approximately 2s = 16 data points for 32s sequence). The response function produced, simulated the BOLD response during the localiser run i.e. when stimuli were being presented in one quadrant per block. This was repeated for each quadrant, and then put together to confirm that the simulated ‘timeseries’ were uncorrelated (see Fig 4.2a).

The correlation coefficient between the individual simulated ‘timeseries’ from the localiser session and the combined simulated ‘timeseries’ for the BCI sessions were calculated. The weighting of the contribution of one sequence (i.e. the 'attended sequence') to the total response was increased in small steps. These ‘weights’ were normalised, and acted to model the effect of attention on one of the quadrant-related timeseries in a BCI session. The correlation coefficient between the individual quadrant-specific timeseries from the localiser session, and the same quadrant in the presence of combined timeseries from the BCI session were calculated. The

higher the correlation for a specific quadrant between the localiser session and the BCI session, the more separable the neural activity linked to the allocation of attention to that quadrant in the presence of competing stimulus streams. I performed this sequence one hundred times for each weighting level with the addition of random noise. The average response frequency was then calculated. Fig 4.2b illustrates the modelled BOLD responses for each quadrant and the effect of 'attention' (i.e. increased weighting on a specific quadrant). This confirmed that the modelled BOLD activity for each quadrant could be distinguished as being different from the other three, motivating the choice of each of the 4 quadrant specific m-sequences.

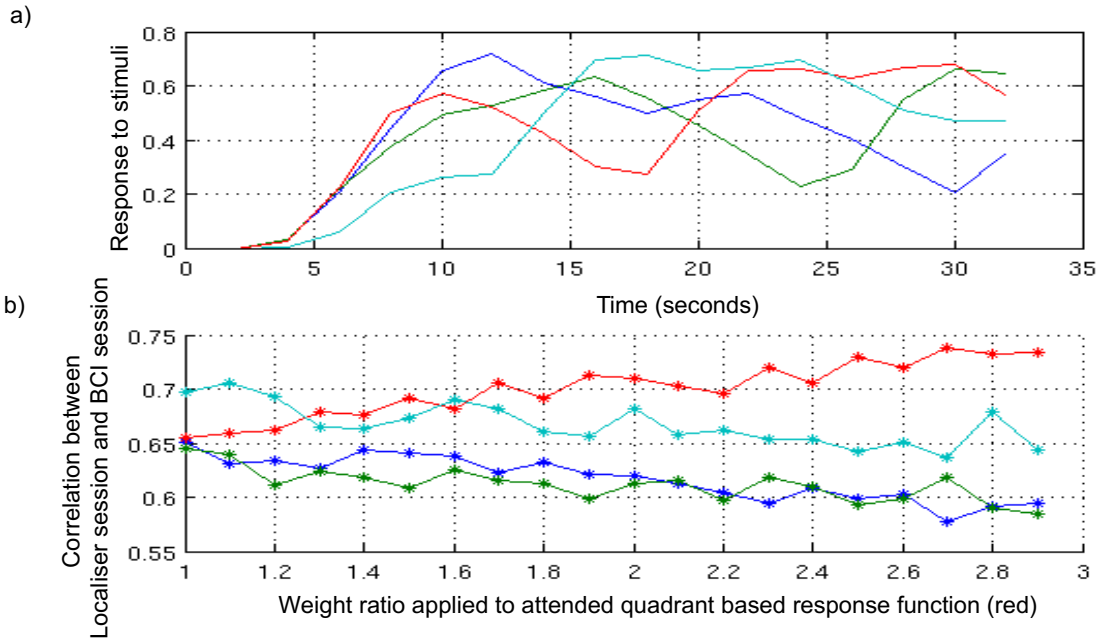

Fig 4-2. Graphs modelling M-sequences for each quadrant convolved with a HRF.

Fig 4.2 a) This shows a timeseries for each quadrant. This figure illustrates the relative orthogonality for each quadrant as represented by a unique and maximally orthogonalised m-sequence convolved with a canonical HRF.

Fig 4.2 b) This illustrates the degree of correlation between the timeseries from the localiser session in which there was only one quadrant based stimulus stream, and the 'BCI' sessions, in which there were 4 simultaneous quadrant based stimulus streams. The objective of this computational model was to examine if the introduction of weighting to one of the time series would make it more discrete from the other three i.e. increase its degree of correlation with the localiser session in which neural activity can only be related to the one quadrant stimulus stream available. The introduction of a weighting served to mimic the effect of attention.

4.2.10 Data acquisition

FMRI data were acquired on a 3T Allegra scanner. Structural images were acquired using a T1-weighted sequence (176 sagittal slices, TR=7.92 ms, TE=2.4 ms, FoV=256x240 mm). Functional images were acquired with a gradient-echo EPI sequence (TR=60 ms, TE=30 ms, FoV=64x72, matrix=64×64, interleaved acquisition, no gap, 3 mm isotropic resolution, 33 slices). Slices were positioned along the slope of the temporal lobe and covered ventral visual cortex and parietal cortex. The sessions of the experiment consisted of 197 volumes each. Field maps were acquired to correct geometric distortions in the EPI images due to inhomogeneities in the magnetic field.

4.2.11 Data analysis- Eye tracking

Eye tracking data were analysed with in-house scripts utilising MATLAB and SPM8. After eye blinks and periods of signal loss were removed from the data, mean and median x and y eye position coordinates were expressed as a distance from central fixation. This was performed for each trial type and for all 8 participants. These data were submitted to a repeated measures ANOVA with the factors of horizontal eye position (left, right) and vertical eye position (upper, lower), to calculate whether mean eye position deviated significantly from fixation across attention conditions.

4.2.12 Data analysis- fMRI data

Participant specific functional data were first analysed using SPM8. To allow for T1 equilibration, the first five images of each run were discarded. Preprocessing of the data involved mean bias correction, realignment (of each scan to the first scan of each run), unwarping, and co-registration of the functional data to the structural scan. A 6mm FWHM Gaussian kernel was used for smoothing the data. The data was filtered with a 256-s cut-off, high-pass filter to remove low-frequency noise including differences between runs, while at the same time preserving as many of the spontaneous fMRI fluctuations as possible (433). The use of the longer length high pass filter was a compromise necessitated by the loss of task related signal with

shorter high-pass filters i.e. 128s. I used session specific grand mean scaling, with no global normalisation when preparing the GLM. Movement parameters in the three directions of motion and three degrees of rotation were included as confounds in the analysis of the imaging data.

4.2.13 Data analysis- retinotopy

Functional data were preprocessed in SPM8 by applying slice time correction, realignment and unwarping (using field maps), and co-registration to the structural scan for each participant. The time series from each scan was then analysed further in MATLAB using a fast Fourier transform and the phase and power at the stimulation frequency (12 cycles/scan for polar scan, 20 cycles/scan for eccentricity scan) were extracted. A F-statistic indicating the significance of the visual response was calculated by dividing the power at the fundamental frequency of the stimulus by the average power across all frequencies. The resulting phase maps were displayed on a reconstructed, inflated surface of the grey/white matter boundary in FreeSurfer (285,292). The boundaries of the visual regions were delineated manually by mapping mirror reversals in the phase map, which correspond to the representation of the vertical and horizontal meridians. The inner and outer boundaries of the ROIs were defined by mapping the extent of significant ($p < 0.05$, uncorrected) visual responses in the polar map. The eccentricity map was used to confirm that these edges corresponded to the locations where the eccentricity values wrapped around. The surface area of each region was then determined by summing up the areas for all the mesh vertices in the region (as calculated by FreeSurfer's reconstruction algorithm for the grey-white matter surface)

Further analysis of the fMRI data will be described in two parts. First, I will explain the analysis of the functional localizer run, the selection of functional ROIs, and then the analysis of the main experiment.

4.2.14 Defining functional regions of interest

For each participant, the parameter estimates resulting from the fit of the GLM to the fMRI time series from the first session were used to calculate a t-statistic at each voxel. This resulted in participant-specific statistical maps of t-values summarising activations associated with the task (in this case summed across all conditions) compared to rest. The t-maps (threshold = 3.0) were overlaid onto each inflated cortical hemisphere (reconstructed from the anatomical images) using Freesurfer (http://surfer.nmr.mgh.harvard.edu/fswiki) (285,292). Participant-specific functional ROIs were delineated manually, including parietal cortex (superior and inferior), fusiform temporal cortex, and LOC.

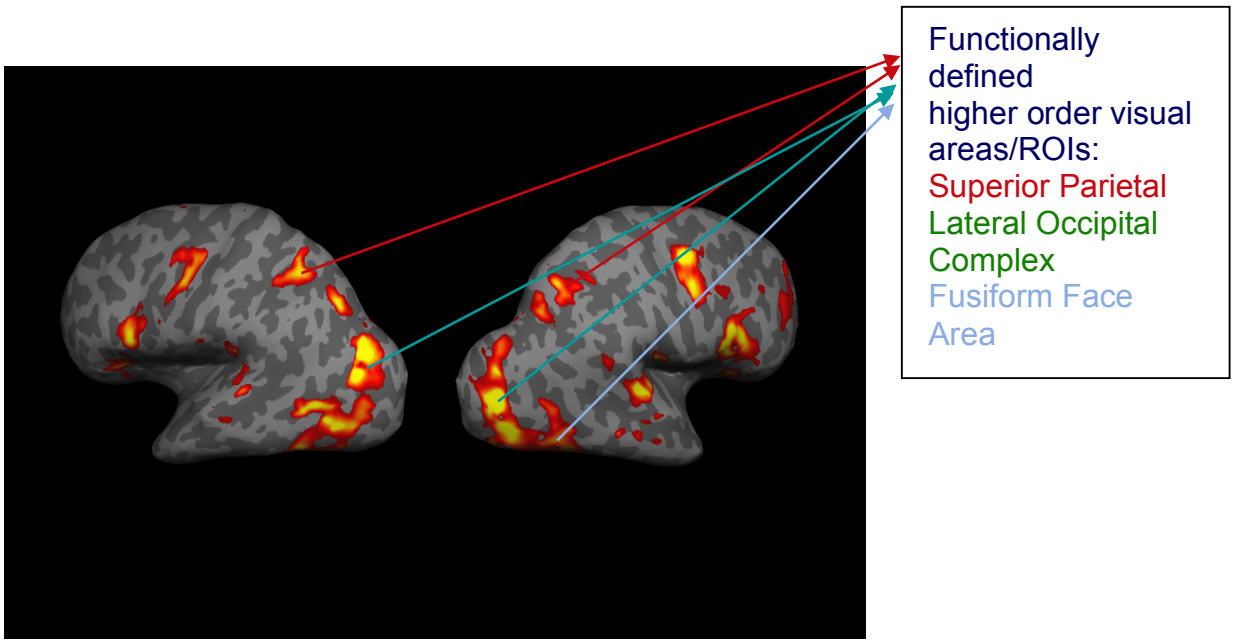

Fig 4-3. Inflated left and right hemispheric brains for a specific participant reconstructed from their anatomical image, showing functional ROIs.

A t-threshold of 3 was used to producing heat-maps of functional activations in relation to the attention localiser task. Using the anatomical labels provided in FreeSurfer as a guide, the resulting overlays were then appropriately demarcated and labelled as isolated participant-specific functional ROIs.

4.2.15 Analysis of main experiment (Sessions 2-5)

Cortical responses to the 4 attentional conditions were modelled using onsets related to the onset times of images within a given m-sequence, convolved with the haemodynamic response. Each m-sequence was specific to 1 of the 4 quadrants; the same m-sequence for a given quadrant was used across all sessions, irrespective of the object category. A GLM was created with separate partitions for sessions 2-5; a regressor was created to model each of the quadrant-specific m-sequences, rest, cue and button box presses, for every block in every session. This yielded a parameter estimate for each condition and for every session.

Establishing which quadrant was attended to during a specific block was based on the highest representative parameter estimate. After establishing this for each block in each session, I then compared the 'predicted' parameter estimate and its quadrant, with the 'chosen' quadrant as indicated by the participant. This was performed on a trial-by-trial basis, enabling prediction accuracies to be calculated across blocks and sessions for each ROI.

4.2.16 Multivariate pattern analysis with 'mini-blocks'

Preprocessed functional data in volume space were further analysed using custom software written in MATLAB. The time course from each run was normalised to a z-score. Four features were extracted from each of the ROIs for each condition (corresponding to the four attention conditions), representing the mean, maximum and minimum parameter estimates together with their standard deviations. Prior to this the time series was modelled in a GLM delineating 'mini-blocks' representing the object category being shown (two categories being shown during the course of a block within any given session of the experiment). For each ROI the features in each mini-block were extracted and vectorised, resulting in one pattern ('sample') for each mini-block.

This data was then used for multivariate classification using a leave-one-run-out cross-validation procedure. Samples from all except one session were assigned to a

training set and the remaining samples were used as a test set. For each attentional condition I calculated the mean sample across all blocks in the training set. To decode I used separate classifiers incorporating each sample in the test set and the mean samples from the training set. All voxels within a functionally defined ROI were used to define features.

4.2.17 Defining Primary Visual Areas (V1, V2, V3)

Primary visual areas were defined using a standard retinotopic approach (284). Specifically, for each participant, using the anatomical images, the surface of each cortical hemisphere was reconstructed and inflated using FreeSurfer. Polar maps of the visual cortex were calculated using phase-encoded retinotopic mapping techniques (284) and retinotopic visual areas were delineated manually. The boundaries of V1–V3 were delineated by identifying the representation of the vertical and horizontal meridians from the mirror reversals in the phase map, separating the ventral and dorsal sub-regions of these areas.

The functional data from the first retinotopic localiser was analysed with a GLM. The effect of activation versus rest was modelled together with movement parameters, and convolved with the HRF. The resulting t-maps were overlaid onto the inflated cortical hemispheres. Twelve discrete regions were uniformly identified in primary visual cortex in all participants, i.e. 3 regions either side of the calcarine sulcus for each hemisphere corresponding to V1v, V2v, V3v and V1d, V2d, V3d respectively. The demarcations between the component regions were delineated using the previously defined retinotopic maps (see above).

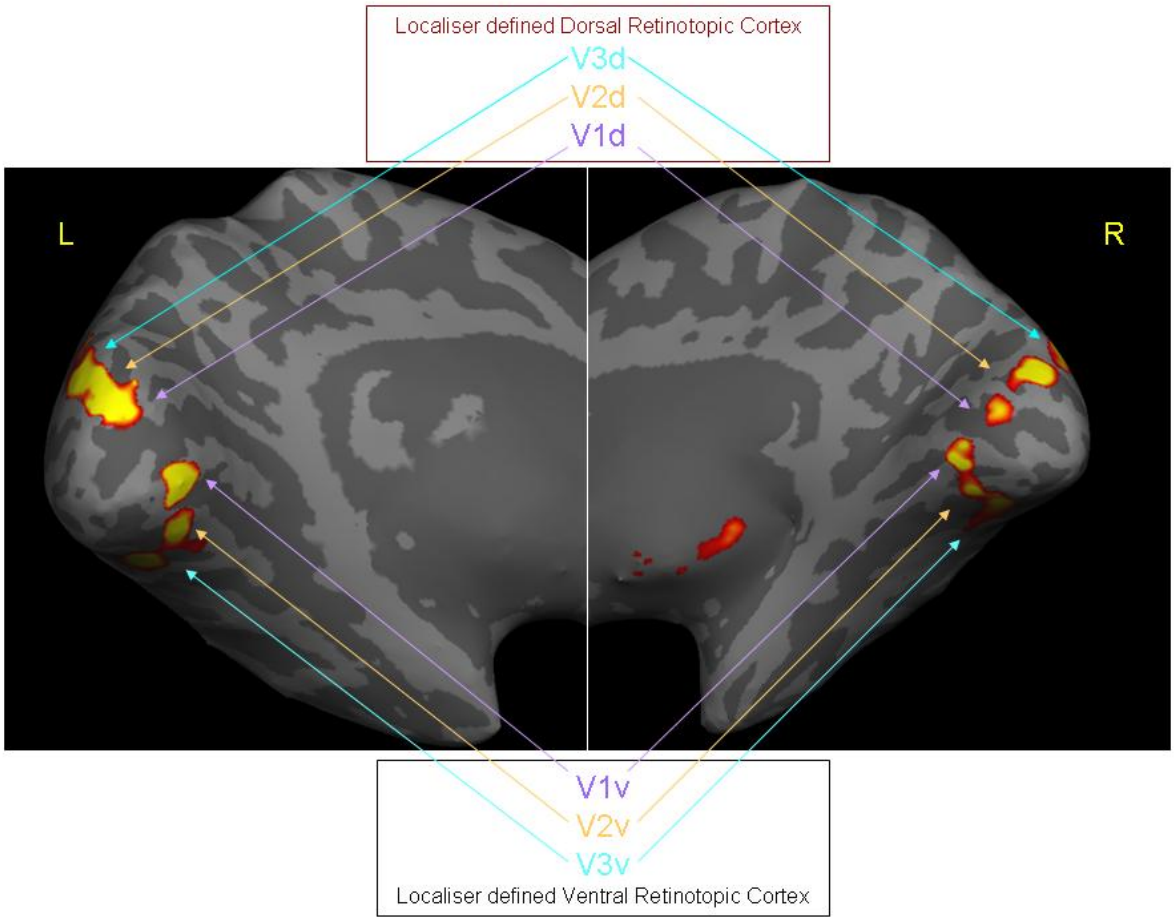

Fig 4-4. Inflated right and left hemispheric brains for a specific participant reconstructed from the anatomical images, showing retinotopic ROIs.

Each hemisphere has been transposed to optimally expose the medial/inferior cortical surface and therefore primary visual cortex. The t-maps associated with the retinotopic localiser taken from the GLM created in SPM, were imported and overlaid onto each appropriate hemisphere. These overlays were thresholded (t-threshold = 3) and smoothed (6mm FWHM kernel) creating heat maps of functional activations, prior to being delineated into each of the specific regions of primary retinotopic cortex. This was aided by the application of overlays of polar and eccentricity maps obtained separately after a phase-encoded retinotopy session

4.2.18 Analysis using retinotopic ROIs

In a similar manner to the main experiment, using each individually defined retinotopic ROI, the parameter estimates relating to the quadrants (modelled using the m-sequence) were examined on a block-by-block basis. This established which quadrant produced the largest parameter estimate/block, providing a prediction of which quadrant had been attended to, which was then compared to the actual attended quadrant.

Each of the 4 component parts of each primary visual area related to one of the attended quadrants i.e. retinotopic cortex which is contralateral and flipped with regards to the viewed real world quadrant (e.g. RV1d relates to lower left visual quadrant). In this manner a competition was set up within each primary visual area during each attentional block, to examine which 1 of the 4 component ROIs had the 'winning' parameter estimate (based on the modelled m-sequences).

I used several approaches for identifying which of the four ROIs within a specific region of primary visual cortex was the 'winner' with regards to a particular attentional task. This involved tailored use of parameter estimates:

4.2.19 Using higher order ROIs:

M-sequence regressor within a ROI

Each m-sequence was specific to a particular quadrant during the experiment. As such within each ROI, there were 4 m-sequence regressors each specific to a complementary area of attended space. I set up a competition among the 4 m-sequence specific regressors, with the highest mean parameter estimate identifying the 'winning' regressor. This analysis examined the classification from within a specific ROI, and the rate of accurate prediction regarding a participant's deployment of spatial attention, on a single trial basis.

4.2.20 Using retinotopic ROIs:

i) Winning parameter estimate

The highest parameter estimate within a ROI (constituting 4 ‘mini’ ROIs within a specific retinotopic region) was used to predict which quadrant of space was being attended.

ii) Mean parameter estimate

In this analysis the four parameter estimates (modelling each of the quadrant specific m-sequences) within a ROI, were established. A competition was then setup among the 4 ROIs, constituting a visual region), with the highest ‘mean’ parameter estimate identifying the ‘winning’ ROI and the area attended.

iii) Median parameter estimate

The approach used a similar principle to ii), using instead the highest median parameter estimate value to identify which of the component ROIs within a specific area of visual cortex provided the highest classification accuracy.

4.3. Results

4.3.5 Eye position data

Participants were requested to maintain central eye fixation throughout all of the 8 task blocks, for each of the 4 sessions. A repeated measures ANOVA was performed on the X and Y eye position data separately, and the factors of horizontal attention (left, right) and vertical attention (up, down) demonstrated no main effect of horizontal or vertical attention and no interaction between them: for X-position data: left vs. right, $F(1,7)=0.697$, $p=0.431$; up vs. down, $F(1.,7)=0.387$, $p=0.554$, interaction, $F(1,7)=1.164$, $p=0.316$; for Y–position data: left vs. right, $F(1,7)=0.697$, $p=0.431$, up vs. down, $F(1,7)=0.387$, $p=0.554$, interaction, $F(1,7)=1.164$, $p=0.316$). Participants therefore did not move their eyes in a statistically significant, consistent manner throughout the experiment, precluding the requirement for the removal of data and/or participants.

4.3.1 Classification within Functional ROIs

I investigated the extent to which functionally delineated high order visual areas could be used to predict the direction of spatial attention. The added caveat was the incorporation of unique temporal information in the presentation of stimuli at each of the 4 quadrant-based spatial locations.

The resulting accuracies for individual ROI based classifications were determined by the extent to which the winning m-sequence parameter estimate, and the linked ‘predicted’ quadrant space, matched the actual attended quadrant during a task block.

Accuracies were calculated using in-house MATLAB scripts, across blocks for each session, across the sessions for each participant, and across participants. The higher order visual areas were grouped into parietal regions, FFA and LOC.

First, I examined the accuracy of classification for the cued session (session 2) across participants (Figure 4-1). Mean accuracies were particularly high in LOC, left inferior parietal ROIs and FFA (See Table 4-1, and Fig. 4-5).

I then examined accuracy of classification for all sessions i.e. 2-5. Classification accuracies were significantly greater than chance in all three higher order visual areas, exceeding 50% in bilateral LOC ROIs, with chance at 25% (See Table 4-2, and Figure 4-6).

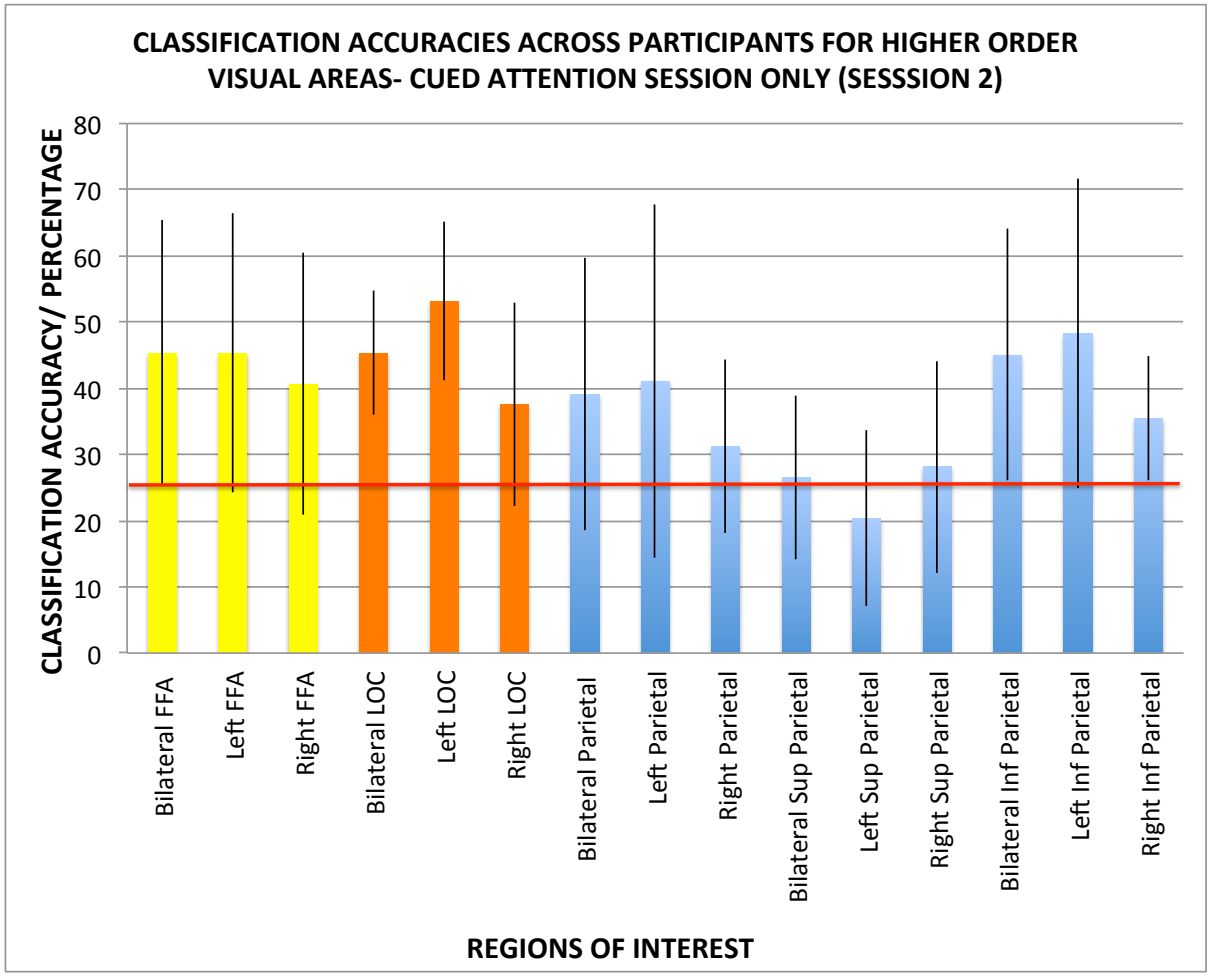

Fig 4-5. Mean classification accuracy for ROIs, across Session 2 (cued attention) and across participants. Accuracy was by calculated by comparing the predicted quadrant with the cue direction on a block-by-block basis. The red line indicates 'chance' level classification at 25%. The error bars indicate ±1 SD.

Table 4-1. Univariate classification accuracies across ROIs, averaged across participants for session 2 (cued attention).

Reported are mean classification accuracies and associated t-values, p-values SD, and degrees of freedom for pairwise t-tests comparing means against chance level classification (25%). Significant classification is marked in red. (Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom).

ROIs	Acc	SD	DoF	t- value	p-value
Bilateral FFA	45	19.97	7	2.8762	0.012
Left FFA	45	21.06	7	2.7277	0.015
Right FFA	41	19.76	7	2.2361	0.030
Bilateral LOC	45.	9.30	7	6.1775	0.000
Left LOC	53	11.97	3	4.7001	0.009
Right LOC	38	15.31	4	1.8257	0.071
Bilateral Parietal	39	20.53	7	1.9378	0.047
Left Parietal	41	26.73	6	1.591	0.081
Right Parietal	31	13.11	5	1.1677	0.147
Bilateral Superior Parietal	27	12.39	7	0.3568	0.366
Left Superior Parietal	20	13.26	7	-1	0.825
Right Superior Parietal	28	16.02	7	0.5517	0.300
Bilateral inferior Parietal	45	18.96	4	2.3591	0.039
Left Inferior Parietal	48	23.31	6	2.6354	0.019
Right Inferior Parietal	35	9.41	5	2.7116	0.021

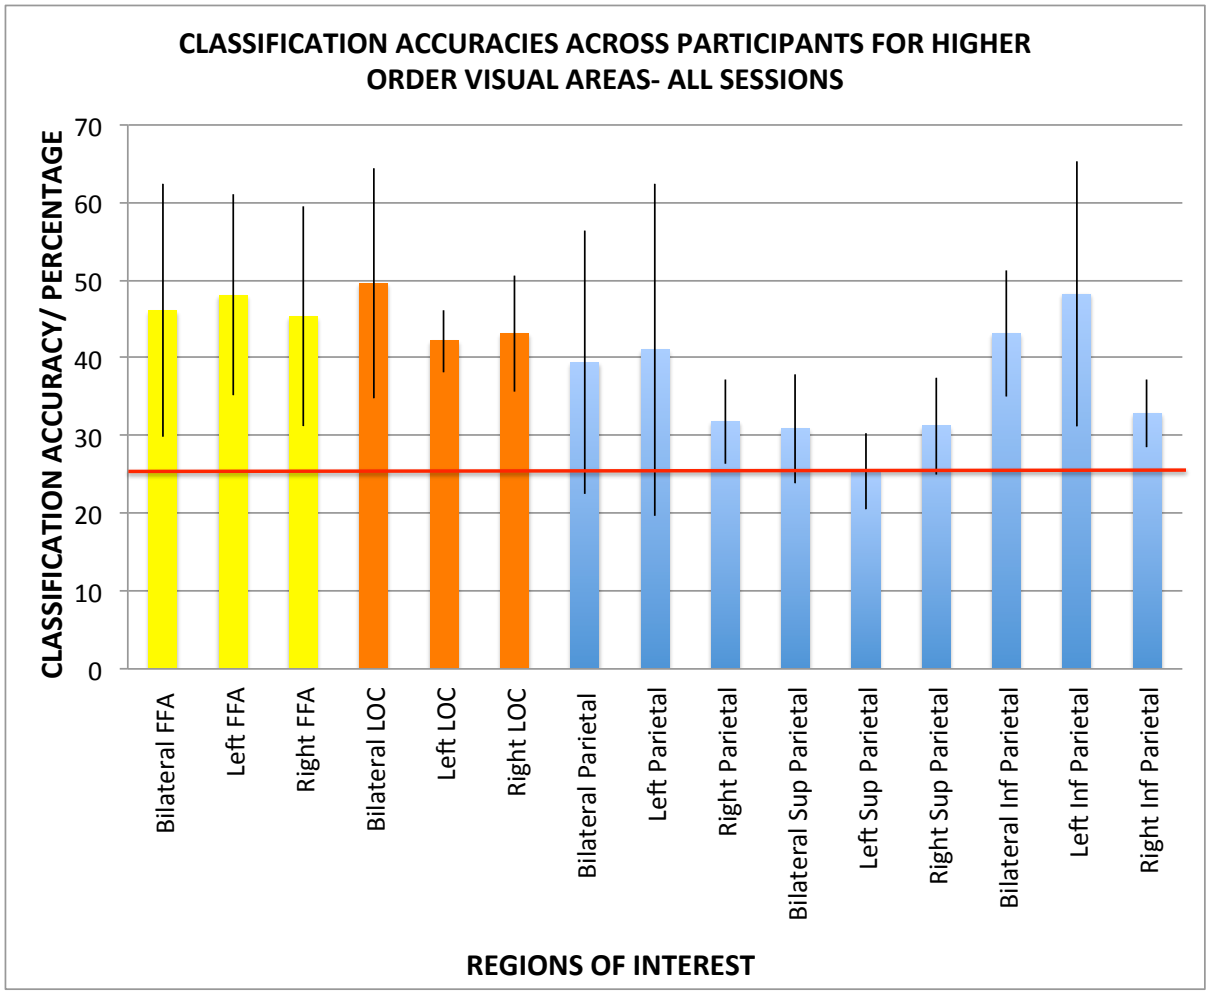

Fig 4-6. Mean classification accuracy for ROIs, across sessions 2-5 and across participants. Classification was based on a winner-takes-all approach for the winning mean parameter estimate. Accuracy was calculated by comparing the predicted quadrant with the attended quadrant (either cued as in Session 2 or selected by the participant Sessions 3-5). The red line indicates chance level classification at 25 %. The error bars indicate ±1 SD.

Table 4-2. Table of univariate classification across ROIs, averaged across participants, for all sessions.

Reported are mean classification accuracies and associated t-values, p- value standard deviation, and degrees of freedom for pairwise t-tests comparing means against chance level classification (25%). Significant classification is marked in red. (Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom).

ROIs	Acc	SD	DoF	t -value	p-value
Bilateral FFA	47	16.26	7	3.67	0.004
Left FFA	48	12.93	7	5.04	0.001
Right FFA	45	14.08	7	4.08	0.002
Bilateral LOC	50	14.80	7	4.71	0.001
Left LOC	42	4.03	3	8.52	0.002
Right LOC	43	7.46	4	5.43	0.003
Bilateral Parietal	39	16.95	7	2.41	0.023
Left Parietal	41	21.38	6	1.99	0.047
Right Parietal	32	5.38	5	3.08	0.014
Bilateral Superior Parietal	31	6.98	7	2.38	0.025
Left Superior Parietal	25	4.85	7	0.23	0.413
Right Superior Parietal	31	6.25	7	2.83	0.013
Bilateral Inferior Parietal	43	8.09	4	5.01	0.004
Left Inferior Parietal	48	17.10	6	3.59	0.006
Right Inferior Parietal	33	4.31	5	4.44	0.003

4.3.2 Classification within retinotopic region using competition amongst constituent ROIs

I next examined the use of primary visual ROIs for classification of the deployment of spatial attention using a univariate approach of identifying parameter estimates acting as markers for the region of attended space. The parameter estimates modelled quadrants as specified by m-sequences.

In the first approach, I utilised a summary value of the activity within a component ROI, in this case the median parameter estimate value. This demonstrated a progressive increase in classification accuracy from V1 to V3, reaching statistical significance at the group level in V3 ($p < 0.05$, for a binomial distribution).

In the second approach I used a competition among the specific parameter estimates representing one of the four component ROIs, modelling the m-sequence for the appropriate region of visual space. This demonstrated increasing above-chance classification through V1 and V2, with a drop of classification accuracy to below chance in V3.

Each of the two approaches used are depicted in Figures 4-8, 4-9 and 4-10, 4-11 respectively. Figures 4-8 and 4-10 show individual participant classification accuracies across each of the three visual ROIs (V1 blue, V2 red, V3 green); Figures 4-9 and 4-11 show group-averaged classification accuracy across each of these three visual ROIs.

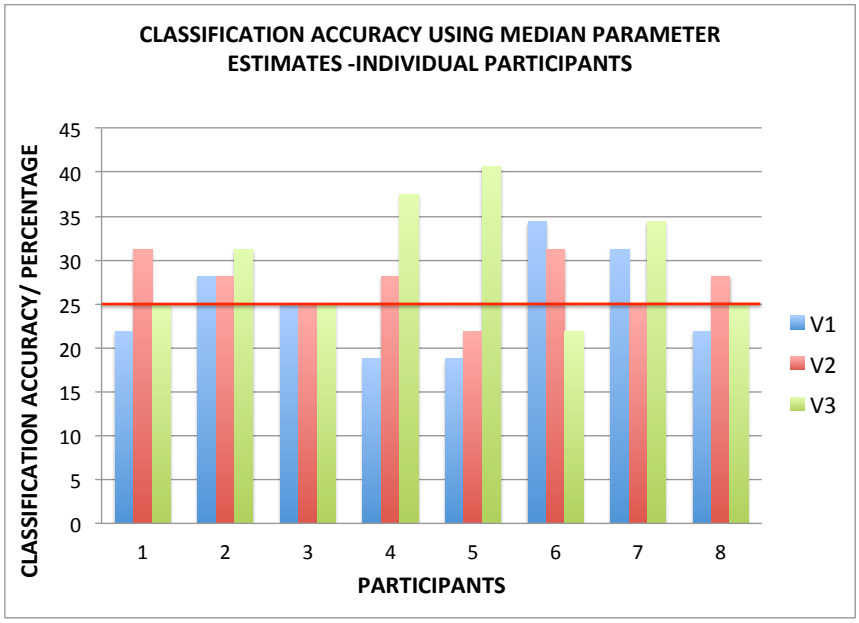

Fig 4-8. Individual classification accuracy using the highest value median parameter estimate to identify the winning ROI, within 3 visual areas (V1-V3). (Chance = 25% indicated by the red line).

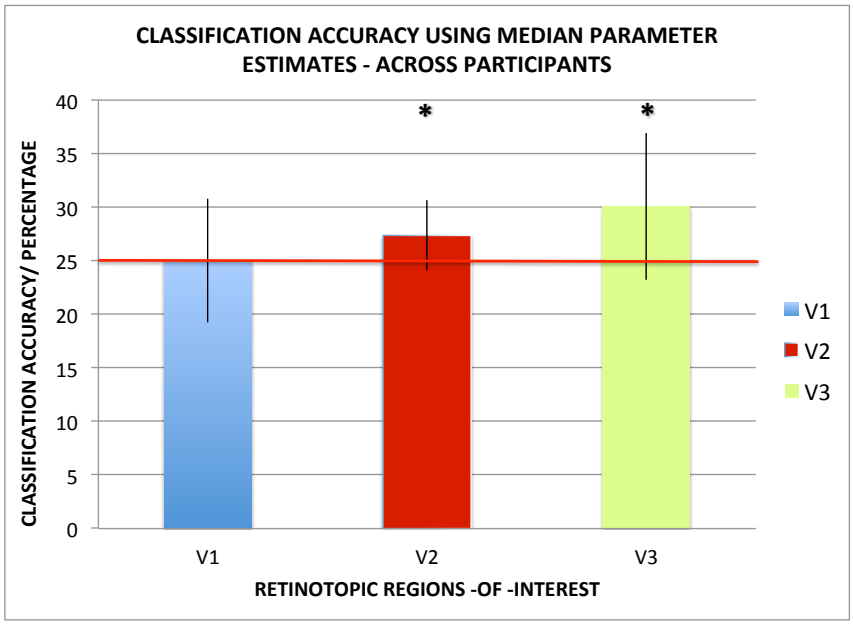

Fig 4-9. Group-averaged classification accuracy using the highest value median parameter estimate to identify the winning ROI, within 3 visual areas (V1-V3). The stars above the bars denote statistical significance ($p<0.05$, one tailed t-test against chance). The error bars indicate ±1 SD. (Chance = 25% indicated by the red line).

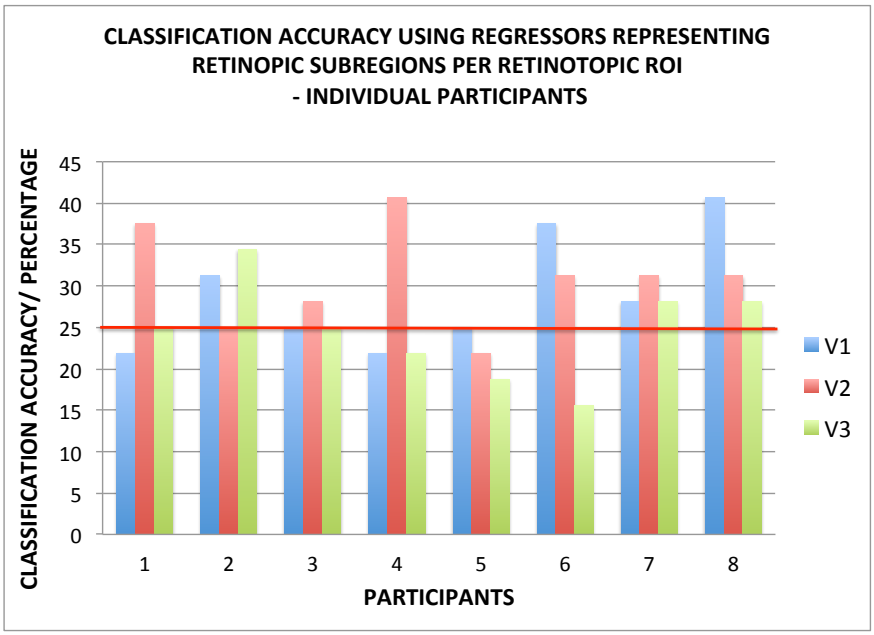

Fig 4-10. Individual classification accuracies based on competition between quadrant-specific regressor taken within a visual ROI. (Chance = 25% indicated by the red line).

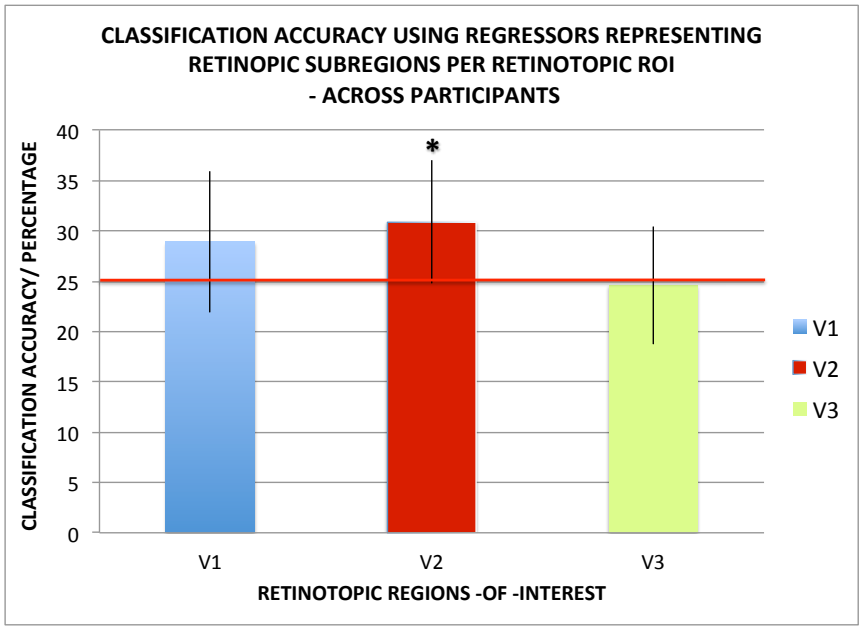

Fig 4-11. Group-averaged classification accuracies based on competition between quadrant-specific regressor taken within a visual ROI. The star above the bar denotes statistical significance ($p<0.05$, one tailed t-test against chance. Error bars are standard deviation. (Chance = 25% indicated by the red line). The error bars indicate ±1 SD.

Table 4-3. Group averaged classification accuracies using the highest value median parameter estimate to identify the winning ROI, within 3 visual areas (V1-V3).

Reported are mean classification accuracies and associated t-values, p- value standard deviation, and degrees of freedom for pairwise t-tests comparing means against chance level classification (25%). (Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom).

ROIs	Acc	SD	DoF	t -value	p-value
V1	22	5.79	7	0.71	0.5
V2	28	3.23	7	2.53	0.039
V3	30	6.87	7	2.55	0.038

Table 4-4. Group-averaged classification accuracies based on competition between quadrant-specific regressor taken within a visual ROI.

Reported are mean classification accuracies and associated t-values, p- value standard deviation, and degrees of freedom for pairwise t-tests comparing means against chance level classification (25%). (Acc = classification accuracy/ percentage, SD = standard deviation, DoF = degrees of freedom).

ROIs	Acc	SD	DoF	t -value	p-value
V1	29	7.04	7	2.05	0.080
V2	31	6.12	7	3.20	0.015
V3	25	5.89	7	0.84	0.428

4.3.3 Classification using multivariate pattern Analysis – Part 1

I also explored multivariate classification, extending the use in Experiment 1 to include more complex classifiers. Initially, I used two linear classifiers (Correlational and LDA), which I had applied successfully in Experiment 1 (Chapter 3). The use of dimensionality reduction was explored with both classifiers. The rationale for the use of a correlational classifier and LDA classifier has been previously discussed (Chapter 3), and is predicated on its simplicity, ease of understanding, reproducibility and comparatively low risk of over-fitting (269).

Table 4-5. Classification accuracy per functionally defined ROI, across all participants in percent. Chance classification was at 25%.

	Left FFA	Right FFA	Bilateral FFA	Right LOC	Left LOC	Bilateral LOC	Left Parietal	Right Parietal
Pattern Correlation	28.32	29.49	25.39	31.64	28.71	28.71	24.61	23.83
Pattern Correlation (with PCA)	28.13	26.37	24.41	29.49	30.08	30.27	23.24	26.37
LDA (with PCA)	29.1	30.66	26.37	30.66	30.27	25.39	26.37	25.78

Classification was above chance in LOC and FFA ROIs, given that chance level classification for 4-quadrant spatial attention was at 25% (Table 4-5). The second table (Table 4-6) identifies the number of participants, out of a total of nine, for whom classification accuracy was significantly above chance. Using lateralised ROIs, FFA classification had as many as 4 participants reaching statistical significance ($p < 0.05$ for a binomial distribution, for LDA with PCA classification) with the right FFA being the most consistent across classifiers (4 participants with Correlation Classification, 3 participants with LDA). The bilateral LOC ROI demonstrated better classification results for the Correlational Classifiers (2 participants without PCA, 3 participants

with PCA). Interestingly bilateral ROIs did not demonstrate an additive effect in terms of classification accuracy, most likely due to the inclusion of non-contributory voxels, i.e. FFA, and potentially more informative voxels being used for classification in the case of the LOC. The parietal ROIs performed less well, with classification not exceeding chance.

Table 4-6. Total number of participants, out of 9 participants, with significant classification per ROI for each classifier. Statistical significance was assessed using a binomial distribution ($p<0.05$ for a binomial distribution).

	Left FFA	Right FFA	Bilateral FFA	Right LOC	Left LOC	Bilateral LOC	Left Parietal	Right Parietal
Pattern Correlation	1	4	1	1	1	2	0	0
Pattern Correlation (with PCA)	0	0	0	2	2	3	0	1
LDA (with PCA)	4	3	0	2	1	1	1	1

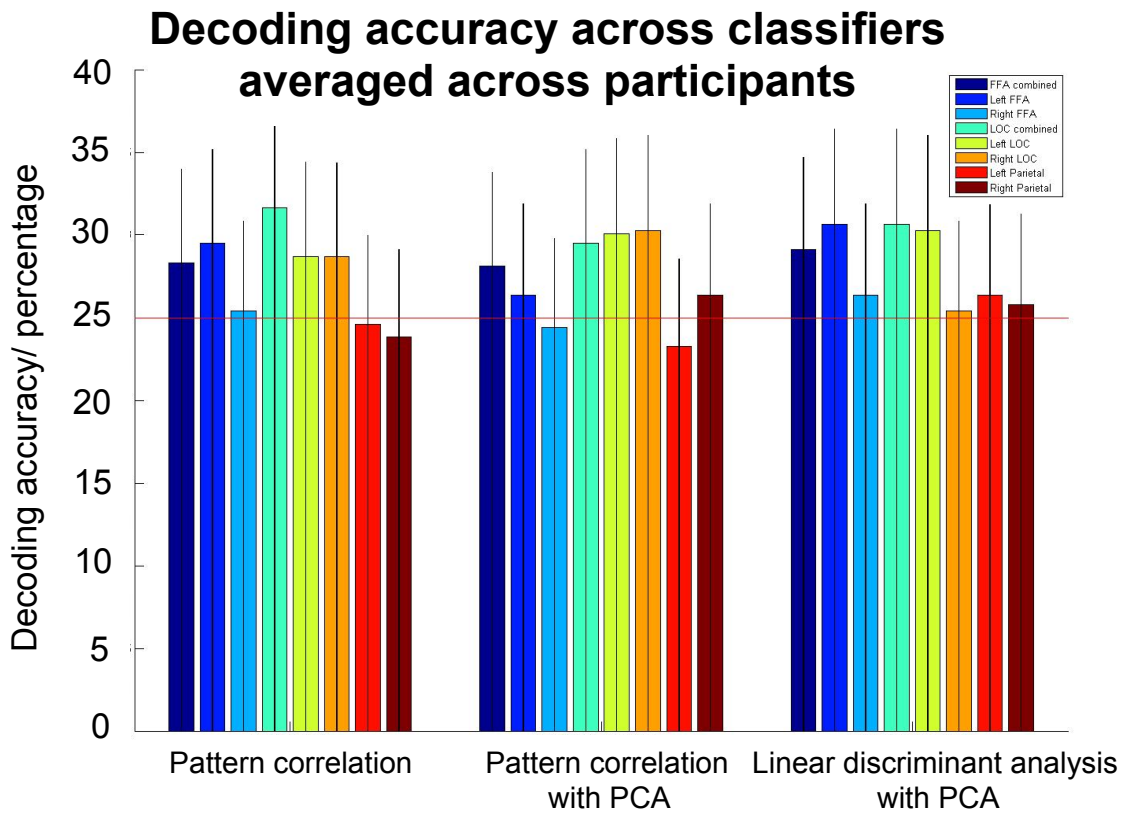

Fig 4-12. Classification accuracies averaged across participants, for each high order visual ROIs, for three linear classifiers. **Chance = 25%, shown by the horizontal red line. The error bars indicate ±1 SD.**

Table 4-8. Table of multivariate classification using LDA across ROIs.

Classification accuracies and standard deviations are shown for most successful linear classifier (LDA with PCA). The number of participants with above-chance classification is reported (p<0.05 for a binomial distribution). (Acc = classification accuracy/ percentage, SD = standard deviation).

ROI	Acc	SD	Participants with significant classification
Bilateral FFA	29.1	0.086	4
Left FFA	30.66	0.075	3

Right FFA	26.37	0.05	0
Bilateral LOC	30.66	0.057	2
Left LOC	30.27	0.052	1
Right LOC	25.39	0.046	1
Left Parietal	26.37	0.066	1
Right Parietal	25.78	0.049	1

4.3.4 Classification using Multivariate pattern analysis – Part 2

I then investigated a series of non-linear classifiers including variations of K-nearest neighbour approaches and SVM. These classifiers have more complex decision boundaries but perform well with smaller ROIs (269). I used the following non-linear approaches: K –nearest neighbour, Quadratic, Radial and Polynomial support vector machines. The use of non-linear classifiers with fMRI data is contentious, with support for their use from one study (258) but not from others (257,273). The expressed concern with the use of more complex non-linear classifiers is that the relationships between features and prediction can become difficult to predict (266), and may not be ideal for fMRI classification.

Linear SVM classifiers have been suggested to be comparable to LDA classification and may be preferable when the optimal decision boundary is approximately hyperplanar, but has assumptions of equal and multinormal pattern distributions (as is the case with LDA) (269).

The purpose of this exercise was twofold. In the first instance I wanted to establish the strength of classification in specific ROIs such as the LOC, which for the previous univariate and multivariate classification approaches had been identified as a putative region from which to decode the direction of spatial attention. The second aim was to identify classifiers that demonstrated comparable accuracies, or better, as compared to the LDA and Correlational Classifier. This is important particularly when considering further developments such as an increased feature space, given that classifiers such as SVM's typically work better with larger numbers of features, and the possibility of combining multiple ROIs.

Examining the total number of statistically significant classifications across all ROIs (see Table 4-10), all the non-linear SVM classifiers performed similar to the LDA classifier, with the radial SVM having the largest number of significant classifications across all ROIs used (i.e. see Table 4-10, 14 cf. 13 with LDA).

Table 4-9. Classification accuracy across a series of linear and non-linear classifiers per functionally defined ROI, averaged across participants. Chance classification was at 25%.

	Left FFA	Right FFA	Bilateral FFA	Right LOC	Left LOC	Bilateral LOC	Left Parietal	Right Parietal
Euclidean Nearest Neighbour	31.64	28.32	26.37	23.63	28.52	23.24	29.88	26.17
K-nearest neighbour 3	28.32	30.08	28.52	25.98	28.52	25.2	27.93	25.59
K-nearest neighbour 5	27.54	28.91	26.76	28.32	29.88	26.37	28.13	23.44
K-nearest neighbour 7	27.73	29.49	25.59	28.13	30.47	25.59	26.17	23.63
K-nearest neighbour 9	25	27.54	25.2	26.76	29.49	26.17	27.15	26.56
Linear SVM	27.34	30.66	25.78	29.88	30.66	26.17	24.22	26.95
Quadratic SVM	24.41	27.15	24.22	29.49	32.23	25.59	29.1	25
Polynomial SVM	28.32	29.1	27.34	28.32	32.03	25.2	25.98	26.37
Radial SVM	25.78	29.3	26.76	29.49	28.91	28.13	29.49	24.61

Table 4-10. Total number of participants with significant classification for each ROI and classifier. ($p<0.05$ for a binomial distribution).

	Left FFA	Right FFA	Bilateral FFA	Right LOC	Left LOC	Bilateral LOC	Left Parietal	Right Parietal
Euclidean Nearest Neighbour	3	2	2	1	1	1	1	0
K-nearest neighbour 3	0	2	3	0	2	1	1	0
K-nearest neighbour 5	0	2	1	1	2	1	1	0
K-nearest neighbour 7	0	2	1	3	1	0	2	1
K-nearest neighbour 9	0	2	1	1	2	0	2	1

Linear SVM	1	2	0	1	3	0	1	1
Quadratic SVM	0	3	1	2	3	0	3	0
Polynomial SVM	1	2	0	2	4	0	0	1
Radial SVM	0	2	1	2	3	3	2	1

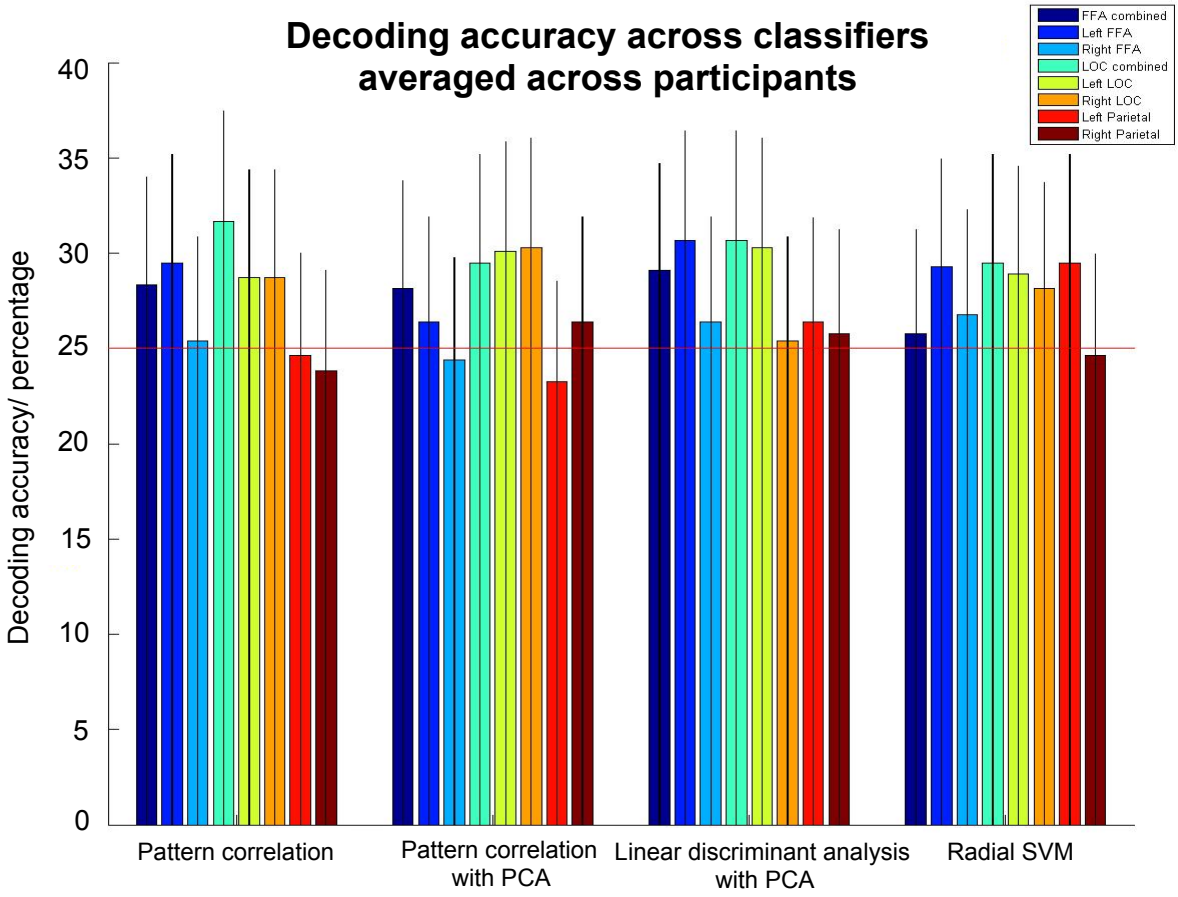

Fig 4-3. Group-averaged classification accuracies, for each high order ROI, for linear classifiers and most successful non-linear classifier.

The most successful of the non-linear classifiers (Radial SVM) is shown in comparison to three previously used linear classifiers. Chance classification is at 25% and is shown by the horizontal red line. The error bars indicate ±1 SD.

4.4 Discussion

In this study, the primary aim was to identify a robust mechanism for classification of covert, participant determined deployment of visuospatial attention using fMRI, on a trial-by-trial basis. Results obtained over all nine participants show that quadrant shifts of spatial attention could be decoded using mean responses with accuracies of approximately 50% (chance 25%) in higher order visual areas including LOC, FFA and in both superior and inferior parietal lobes. I compared univariate and multivariate classification, with the aim of optimising a classification platform for future implementation as a brain-computer interface. I found that univariate classification performed best, using higher-order bilateral ROIs.

The analysis can be considered in three parts. Firstly, classification based on univariate analysis, examining which parameter estimate (modelling each quadrant-specific m-sequence) predicted the direction of covert attention, performed with the highest accuracy. Accuracies approximately twice chance on a single trial basis were obtained in LOC and FFA. Secondly, the use of multivariate classifiers, in particular the LDA and the radial SVM classifiers demonstrated above chance accuracy in similar higher order visual areas including LOC and parietal lobe, although at much lower levels of accuracy.

The multivariate classifiers were applied to blocks of imaging data, specifically smaller ‘mini-blocks’ defining particular object category; feature selection included the mean, minimum and maximum parameter estimate for each quadrant. It is worth considering that despite the spectrum of accuracies achieved across the different classifiers and functional ROIs, persistently high classification accuracy was evident in the LOC. Taking the two analyses together, this suggests that LOC may represent a putative target for extracting object-based and spatial information. As discussed in the previous chapter, it has been reported that the LOC in addition to containing object information, may also contain spatial information. There is evidence to suggest that the object information may be location specific, or location tolerant (48). In this study, prediction accuracy for the direction of spatial attention may in fact be

dependent on the presence of retinotopic maps in LOC. A similar mechanism may underlie the prediction accuracies observed in the inferior temporal cortex e.g. FFA. The existence of retinotopic maps within the LOC (112,384), and the inferior temporal cortex (FFA and PPA) (41,379,410,434) has been previously demonstrated.

I performed formal retinotopy followed by an assessment of the strength of predictive classification in primary retinotopic regions (see Figs 4-8 to 4-11). I observed an incremental trend in classification accuracy using median parameter estimates through V1-V3, with significant classification in V3 only. This would be in keeping with the current understanding of attentional modulation, which increases progressively through primary visual areas, with V1 having little in the way of 'attentionotopy', V2 more so, and V3 demonstrating significant attentionotopy (435). The use of median values enables the best measure of the central location of data in the presence of a highly skewed distribution, and is less susceptible to the influence of outlying values. In this context, the median parameter estimate is likely to be the measure that is the least susceptible to noise in the BOLD signal, and that best represents the activity within a component ROI.

The use of parameter estimates modelling a specific subdivision of the primary visual cortex was predicted to be the best means of classifying attention using retinotopic cortex. This ostensibly sets up a competition between four-component 'mini' ROIs linked to retinotopic cortex serving specific quadrant of visual space. The analysis indicated that V2 might have a slight advantage with regards to prediction accuracy. However, the use of retinotopic ROIs produced lower rates of classification accuracy overall as compared to predictions made using higher visual areas.

The use of quadrant specific timing of stimuli presentations may have contributed to the increased classification accuracies demonstrated in high order visual areas. The use of timing to give distinguishing information to cognitive experiments performed using fMRI originally led to the development of event-related fMRI. As compared to traditional block design fMRI paradigms which combine identical event types in a row, event related fMRI allows for different event types to be interleaved. Therefore, event related fMRI provides an opportunity to resolve haemodynamic responses

linked with neural activity produced by multiple stimuli occurring close together in time, as compared to block design fMRI which is optimised for detecting activity in the brain. In this study, I combined a standard block design with regards to the presentation of object categories (i.e. mini-blocks of faces, house, body parts, everyday objects), with specifically timed presentations for individual stimuli, optimised through the use of m-sequences. By ensuring maximal orthogonality between the 4 quadrant stimulus presentations, the use of m-sequences had the defined purpose of helping to separate brain activation produced by the deployment of spatial attention to stimuli in a specific quadrant.

A consideration in this experiment, as with the previous chapter, was the examination of multiple brain regions, as well as classification approaches, and therefore the potential of problem of multiple comparisons. As previously discussed, a Bonferroni correction can be applied. This attempts to address the family-wise error rate, that is the likelihood of falsely reporting a significant classification (i.e. a Type 1 error), when multiple brain regions are being examined simultaneously. For n comparisons, the correction is made by adjusting the significance by 1/n. Although applying this correction to the classification results would not affect all of significant results presented in this study, Bonferroni corrections in the context of BCI bio-signal processing has been suggested to be too conservative (411). This relates to the fact that it assumes independence of comparisons, however, BCI classification is frequently performed on co-dependent temporal and spatial data. As discussed in the previous chapter, it may instead be reasonable to adopt a hierarchical significance testing approach. This has been used in EEG analysis (419), and may be applied in principle to fMRI analysis. In this case, the data is put into a ‘family hierarchy’, with sub-families (child hypothesis) e.g. a family of parietal ROIs. Child hypotheses are compared recursively if their ‘parents’’ null hypothesis is rejected. This experiment (and Experiment 1) were fundamentally aimed at investigating which brain regions were optimal for classifying spatial attention, and by necessity tested several classifiers and brain regions. (NB: In Experiment 3 (see rationale, next) I hone the approach to the use of three bilateral ROIs as the targets for the application of univariate classification only).

During the design of this experiment, consideration was given to the requirement for real-world applicability both in terms of the ease of translation of the basic experiment into a real-world setting i.e. by using stimuli with environmental valence (cf. scrambled images, Gabor patches) as well as the technical considerations for mutability into a rt-fMRI classification/ BCI. I successfully established an offline classification pipeline utilising a univariate analysis approach. The established classification framework provides a means of decoding neural activity related to the deployment of spatial and object based attention. This now provides the basis for an online fMRI-based, attention-driven communication interface which uses neural activity as a surrogate for communication (15).

4.5 Conclusion

In this experiment I provided evidence of successful classification of the deployment of 4-quadrant spatial attention using univariate and multivariate approaches in the higher order visual regions examined. A comparison of the ROIs, including higher order regions as well as primary retinotopic cortex, identified higher order ROIs as the most successful target for classification and decoding of spatial attention. This 'decoding' pipeline was optimised by the inclusion of temporal, object and spatial information, and performed most successfully with simple univariate analysis and classification.

4.6 Rationale for Experiment 3

In Experiment 2 I expanded on my previous work by exploring a number of options for the development of a non-invasive, attention-driven fMRI based BCI. I showed that by using a novel implementation for the timing of stimulus presentations, a univariate classification algorithm could successfully decode the direction of covertly deployed spatial attention, at more than twice chance, based on BOLD activation in higher order visual brain regions. Classification accuracies were higher than those achieved with the multivariate classifiers. In the next chapter, I will seek to implement

this into an operational BCI framework, using a rt-fMRI pipeline, providing proof-of-principle for an online ‘attention-driven’ BCI for communication.

Abstract: Experiment 3

Real-time decoding of spatial attention in higher order visual areas

Introduction

Higher order visual areas such as parietal cortex, FFA, and LOC carry information regarding both object category and their spatial location (1). Information from areas such as these can be used to decode object categories, or the direction of covert spatial attention (2,3). I use an automated decision criterion to evaluate, on a trial-by-trial basis, which location is covertly attended by participants performing a 4-quadrant attention task with real-world stimuli, and provide proof-of-principle for a rt-fMRI BCI for communication.

Methods

Eight healthy adult volunteers underwent fMRI scanning, during which they performed a task requiring the covert allocation of attention to one of 4 spatial quadrants. During a visual stimulation block, 4 stimuli (faces, houses, body parts, and consumables) were presented simultaneously in each of the 4 visual quadrants. Stimulus presentation timing followed a quadrant-specific m-sequence, in order to optimise the separation of neural activations related to each of the 4 quadrants. Participants were instructed to covertly attend to one of the 4 quadrants and indicated at the end of each block which quadrant they had attended.

Participants' BOLD signal changes were extracted from the target brain regions (bilateral FFA, bilateral LOC, bilateral parietal), delineated initially with a localiser session. Each of the quadrant parameter estimates were modelled with a GLM and ‘decoding' was carried out in a 'winner takes all' approach, based on which one of the 4 quadrant-specific parameter estimates had the highest mean value .The decoded 'prediction' was compared against the participants' subsequently disclosed

response - this was used to generate a measure of online prediction accuracy across blocks and sessions for each ROI.

Results

The decision rule (i.e. winner-take all, quadrant-specific parameter estimates) produced accuracies in the selected ROIs well above chance (> 25%). Accuracies for determining the direction of attention were significantly above chance in all 3 ROIs (all *p's* < 0.001); individual decoding accuracies were up to 70% accurate during the first half of each experimental session.

Conclusions

Determining the direction of spatial attention using a combined univariate approach on a single-block basis can be achieved using rt-fMRI. Decoding accuracies were significantly better than chance in higher order visual areas including bilateral parietal cortex, LOC and FFA. Decoding spatial attention using temporal information (i.e. m-sequences) and real world stimuli provides a novel approach to the development of 'attention based' BCI (4). These 'real-time' BCIs could utilise attention as a means of communication in patients who are unable to speak, or move e.g. locked-in patients.

Keywords:

rt-fMRI, BCI, attention, online decoding, visual areas, FFA, LOC, parietal

5. EXPERIMENT 3

Real-time decoding of spatial attention in higher order visual areas

5.1 Introduction

Online brain activation, linked to the voluntary control of visual attention during fMRI scanning, may be utilised as a surrogate for communication, leading to the possibility of a rt-fMRI-based BCI communication device. Communication interfaces based on non-invasive, and invasive, imaging modalities aim to provide effective information within a constrained timeframe, ideally on a trial-by-trial basis. Using an automated decision criterion to evaluate shifts of spatial attention by participants performing a 4-quadrant attention task, I sought to provide proof-of-principle for an rt-fMRI-based BCI for communication.

Non-invasive 'decoding' of brain activations with rt-fMRI, is based on the online classification of imaging data, following extraction and processing of task-related BOLD activation performed 'on-the fly'. Any brain state may potentially be decoded in this manner, including those associated with covert cognitive brain states (38,436). Potential applications include 'lie-detection' devices, which have already been demonstrated offline, and the online conversion of brain activations into motor and cognitive output. Cortical activations decoded in this manner, may be explicitly related to the task e.g. motor cortex activations used to move a robotic arm (7). Alternatively, higher order cognitive processes linked to more distributed, complex neural activations, may be harnessed and converted to bits of information for a more versatile BCI (9,432,437). This has included spelling devices, and the spatial navigation of a cursor through a virtual maze (39,40) .

Voxel-based activations, which constitute the make-up of an fMRI image, are the basis of information transfer in fMRI-based BCIs. Voxels themselves range in size from e.g. $1mm^3$ to $3mm^3$, and may contain many neuronal sub-populations with differing tuning properties. In order to optimise signal extraction for the purposes of a BCI, it is necessary to maximally target specific neuronal populations within a cortical region, which are functionally involved in the task being used to drive the BCI. A challenge in the selection of cortical regions for this purpose is ensuring specificity. A potential solution to this problem is to target focal regions of the brain, which have a higher order role in co-ordinating neural processes, and may combine associated neural firing. A putative region would still have differing populations of neurons in close proximity, but now potentially with mutual firing properties – this might then serve as an optimal target for a neural interface.

Regions such as the LOC, FFA and regions of the parietal cortex (i.e. superior and inferior parietal lobule) are candidate regions for an integrative role in processing of visual stimuli and attention. LOC is traditionally recognised as being object selective cortex. However it contains at least 2 retinotopic regions, and additionally appears to be involved in coding inter-object relations (7,8,9). FFA is predominantly face selective, but has also been shown to have spatial topography (41,440). It responds more broadly to visually presented shapes (441), and its activity is modulated by attention (442,443). Parietal cortex has a role in top down processes, including mediating shifts of attention, maintenance of saliency maps, expectation, priming, shape discrimination, but also in aspects of working memory (46,81,444,445). More specifically, dorsal regions of the parietal lobe are activated by shifts of top down attention related to spatial location, features or objects (118,446–449).

In progressing towards building a attention-based BCI, these three regions have been shown to participate in spatially specific attentional shifts, and are actively involved in the top-down modulation of attention, serving as a common point for multiple sources of information (450,451). At this level of the visual hierarchy, there may be sufficient quantitative and qualitative information integration, to enable decoding of veridical representations. This is relevant for a BCI using higher order visual processing such as attention, together with real-world stimuli, for the purposes of a communication interface in patient populations.

The use of rt-fMRI as the basis for a non-invasive BCIs is centred on its superior spatial specificity and resolution, as compared to other comparable non-invasive imaging modalities used in this way e.g. MEG/EEG based devices (2). Functional brain regions can be identified, and investigated in relation to task-based activations, and can be used to hone activation profiles, to increase signal strength, and improve decoding accuracy (10). Further, the neurobiology of the region can be explored, enabling more complex assessments such as examining network connectivity between associated functional brain regions. Finally, an important translation of this sequence is the identification of an anatomical substrate for direct implantation of intracortical BCIs, which has been demonstrated (10). Rt-fMRI 'verified' implantable BCIs e.g. (19) provide an optimal means of validating BCI technology prior to surgical implantation, which has an associated morbidity and risk of mortality (20–23), thereby increasing the safety and specificity of the BCI technique, as well as the likelihood of long-term success.

I aimed to test whether brain activations in higher order visual cortex produced by shifts of spatial attention, covert or otherwise, could be accurately classified in real-time. By doing this I intended to provide proof-of-principle for online 'decoding' of visuospatial attention, which might then be applied in a BCI, to indicate spatial preferences to 1 of 4 quadrants. I hypothesised that neural activations linked to the allocation of spatial attention, could be enriched by the inclusion of information related to stimulus being attended (i.e. object and feature-based information) and its temporal profile. I tested this hypothesis by creating quadrant specific m-sequence timings for the presentation of quadrant specific stimuli. Classification was performed online, on brain activations extracted from bilateral FFA, PPA and parietal cortex, and compared visual quadrant based parameter estimates in a winner-takes- all approach. I predicted that using this method, decoding of visuospatial attention would be most accurate using higher order visual areas.

5.2 Methods

5.2.1 Participants

Eight healthy adult volunteers (24–32 years of age; mean age = 28 years, 4 females) with normal or corrected-to-normal visual acuity were recruited to participate in the experiment. Each participant provided written informed consent and the study was approved by the local ethics committee (UCL Ethics Committee code: 09/H0716/14). Each participant also passed MRI safety screening approved by the Wellcome Trust Centre for Neuroimaging (where MRI scanning was performed).

5.2.2 Stimuli

The visual stimuli consisted of 4 categories namely *faces, houses, body parts,* and *food/drink.* Each category contained 16 unique exemplars that were specific to each quadrant. Each stimulus subtended 2 degrees of visual angle in diameter, and were presented at an eccentricity 6 degrees from the centre of the screen. All images were rendered to ensure identical greyscale values, and mean luminance using a custom designed MATLAB script.

5.2.3 FMRI scanning

All experiments were performed on a 3T Allegra head-only scanner, using a standard transmit–receive head coil. Functional data were acquired with a single-shot gradient echo planar imaging sequence (matrix size, 64_64; field of view, 192_192mm; isotropic resolution, 3 **x** 3 **x** 3mm; 32 slices with ascending acquisition; slice thickness, 2 mm; slice gap, 1 mm; echo time (TE), 30 ms; repetition time (TR), 1920 ms; flip angle, 90°; receiver bandwidth, 3551 Hz/pixel). In the middle of each scanning session, double-echo fast, low-angle shot sequence (FLASH) field maps (TE1, 10 ms; TE2, 12.46 ms; resolution, 3 **x** 3 **x** 2 mm; slice gap, 1 mm) were

acquired and used to correct geometric distortions in the images attributable to field inhomogeneities.

5.2.4 Real-time set up

I used the Turbo Brain Voyager package with custom real-time image export tools programmed in ICE VA25 (Weiskopf et al., 2004b), and custom scripts running on MATLAB. Participants' brain activations (differences in BOLD signal) were extracted from previously identified brain regions i.e. target ROIs, with a delay of 2 s from the acquisition of the image. Head motion was corrected in real time using Turbo Brain Voyager.

The real-time data preprocessing was performed in Turbo Brain Voyager and encompassed 3D motion correction with realignment to a preselected template, smoothing, incremental linear detrending of time series and statistical parametric mapping. The ROI time course(s) were extracted from the prescribed ROI masks, averaged and exported by Turbo Brain Voyager. Signal drift, spikes and high frequency noise were further removed in real time from the exported time courses with the custom MATLAB scripts (302).

5.2.5 Experimental procedure

There were 5 sessions per participant. They consisted of the following :-

1. Localiser session

During the first session, streams of stimuli were presented in one quadrant of the screen for the duration of one block (48 seconds made of 3 seconds cue presentation, 32 seconds stimulus presentation, 13 seconds rest), such that each quadrant was attended to twice over the course of a session. There were 8 blocks in this session, which lasted 6 minutes 24 seconds.

This session served to identify object and spatiotemporal selective regions of higher order visual cortex (i.e. functional localiser). Participants were instructed to maintain central eye fixation throughout the session.

Defining functional regions of interest

ROIs were selected using the Turbo Brain Voyager selection package. For each participant, the parameter estimates resulting from the fit of the GLM to the fMRI time series from the first session were used to calculate a t-statistic at each voxel indicating evidence of task-related activation. This resulted in participant-specific statistical maps of t-values summarising activations associated with the task compared to rest. The t-maps were overlaid onto inflated cortical hemispheres using the Turbo Brain Voyager package (threshold of t=3) Participant-specific functional ROIs were delineated manually and included bilateral parietal cortex, FFA and LOC.

2. Cued attention session

In each task block, object images from a particular object category were presented in a pseudo-random manner, consisting of 16 images per mini-block (400ms per image, 100ms gap); each task block was composed of 2 mini-blocks, resulting in all 16 exemplars within an object category being shown during a mini block. The order of the 'object' mini-blocks was counter-balanced between and across sessions. Attention to a particular quadrant was indicated using a directional cue, presented during the cue interval. Each quadrant was cued for attention twice over the course of the session. Participants were instructed to maintain central eye fixation throughout the session. There were 8 blocks (48 seconds each) in the session which lasted 6 minutes 24 seconds in total.

To ensure attention was maintained in all sessions a button press was required if 2 successive exemplars were identical (i.e.one-back task). This occurred between one to 3 times per mini-block. The n-back task ensured focused attention through the presentation of stimuli.

Cued attention

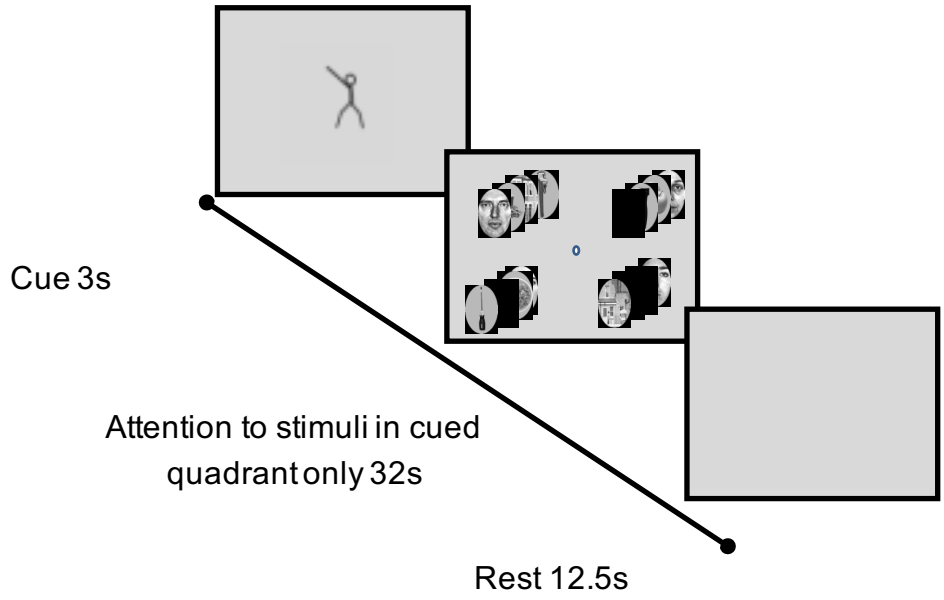

Fig 5-1. Cued Attention session schematic. Participants were cued to attended stimuli presented in one quadrant per block.

4. 'Decoding' sessions

Stimuli were presented as described in the previous paragraph. Participants were now instructed to covertly attend a quadrant of their choice for the duration of a whole block while maintaining central eye fixation. They were further instructed to use a strategy that would enable them to attend all quadrants twice over the course of the scanning session. They disclosed the attended quadrant using a button press during the rest period, at the end of each of block.

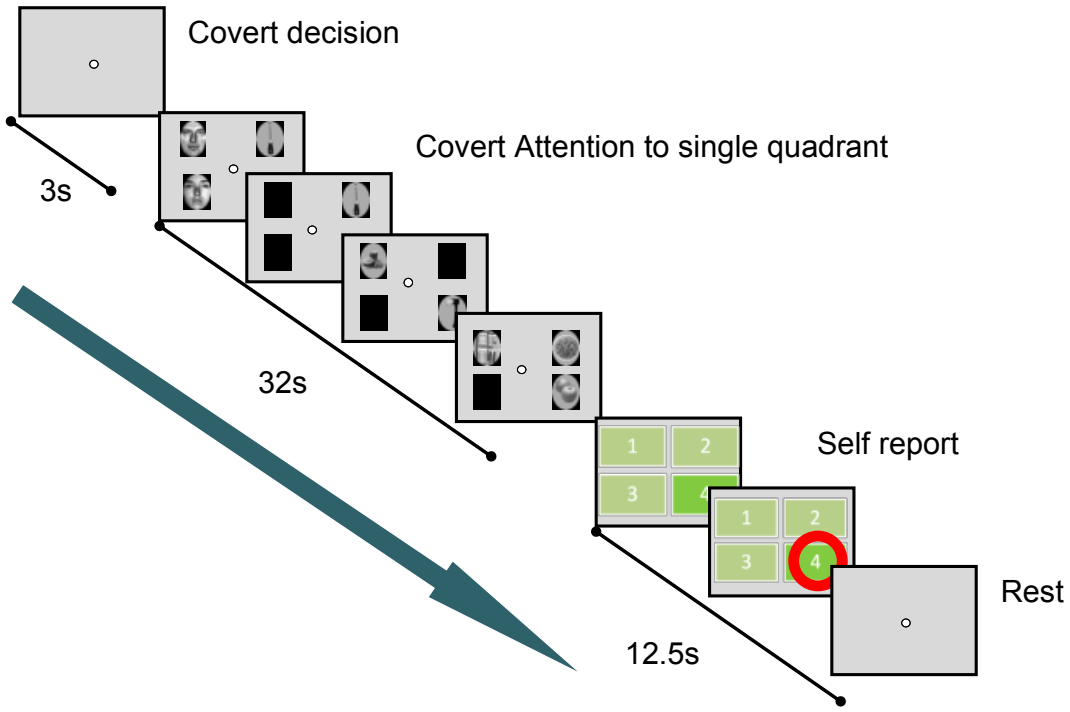

Fig 5-2. 'Decoding' sessions schematic

Participants were instructed to fixate centrally, attend to one of the 4 quadrants during the cue, disclosing which quadrant they had attended at the end of each block using a button-box. Stimulus presentation timing was quadrant specific, specified through the use of m-sequences.

5.2.6 M-sequences

The timing of presentation of the stimuli was prepared using a quadrant specific m-sequence (249). The m-sequences were prepared to ensure maximum orthogonality, providing 32 stimulus presentation slots per quadrant per block. A model of brain activations using maximally orthogonal m-sequences (64 time points of either stimulus/ or blank presentations) was first tested, convolved with modelled 'noise' and the HRF (see Chapter 4).

Frequencies were downsampled to the typical TR (i.e. approximately 2 seconds = 16 data points for 32 second sequence). The sequences were used to simulate the localiser run, with responses to the sequence in one quadrant at time, as would be the case in the session one localiser. The sequences were then used to simulate sessions that had simultaneous streams in all 4 quadrants (i.e. brains activations in response to all of the sequences being shown simultaneously as seen in sessions 2 to 5). The weighting of the contribution of one sequence to the total response (i.e. the 'attended sequence') was varied in small steps. The ‘weights’ were normalised, with weighting representing the amount of attention to a particular sequence. The correlation coefficient between the responses to individual sequences from session one and the combined responses from the multiple stream sessions (2 to 5) were calculated. The sequence was performed one hundred times for each weighting level with the addition of random noise to the responses. This confirmed that the modelled BOLD activity for each quadrant could be distinguished as being different from the other three, motivating the choice of each of the 4 quadrant specific m-sequences.

5.2.7 Eye tracking

Eye tracking was performed during the offline classification study, using a similar experimental set-up (see Experiment 2). This provided support for the absence of excessive eye movements following participant instruction to maintain central eye fixation. I opted not to use eye tracking for the online experiment due to the complexity of the real-time implementation.

5.2.8 Analysis of main experiment (Sessions 2 to 5)

I investigated the extent to which functionally delineated higher order visual ROIs could be used to predict the direction of spatial attention. The addition of unique temporal information in the presentation of stimuli at each of the 4 quadrant spatial locations was applied to improve decoding accuracy. The resulting accuracies for

individual ROI based classifications were based on comparing the winning quadrant specific parameter estimate, with the covertly attended quadrant during a task block.

Cortical responses to the 4 attentional conditions were modelled using onset times of images, together within a given m-sequence, convolved with the HRF. Each m-sequence was unique and specific to one of the 4 quadrants; the same m-sequence for a given quadrant was used across all sessions, irrespective of the object category. A GLM modelled each of the quadrant parameter estimates over each block consisting of 24 scans. 'Decoding' was carried out at the end of each block, in a 'winner takes all' approach, based on one of the 4 parameter estimates with the greatest mean value. Data were read by the script and lagged behind image acquisition by approximately 2 seconds.

The attended quadrant, during a specific block, was the one with the highest representative parameter estimate. On a trial-by-trial basis, a prediction was made, which could then be compared to the actual quadrant attended to by the participant, by comparing the prediction to the participant's button-box response at the end of each block. Having established this for each block in each session, it was possible to calculate accuracies for each ROI.

I previously confirmed offline that the optimal implementation for decoding spatial attention was through the use of higher order visual areas, and applying univariate decoding (see Chapter 4). For this reason, I did not use primary retinotopic ROIs, nor multivariate classification techniques.

5.2.9 Reaction times

The potential effects of experiment time on decoding accuracy are an important consideration for an accurate BCI. A possible effect might be a decrease in decoding accuracy due to increasing fatigue of an individual. I therefore further examined the changes in decoding accuracy over time. I did this by dividing each session, into the first 4 blocks and the second 4 blocks, and comparing reaction times during the

performance of a n-back task, between the first half and the second half of the session. I also did a comparison of the reaction times from one randomised set of four blocks with another, to confirm that any changes observed were not occurring due to chance.

5.3. Results

5.3.1 Decoding accuracies

I first examined decoding accuracies across all sessions and blocks, for each of the 3 bilateral ROIs (See Fig. 5-2), to establish whether the quadrant to which attention was being directed could be decoded at above chance levels from signals evoked in each ROI. For each ROI (FFA, PPA, parietal), decoding accuracy was significantly above chance levels (25%): FFA ($M=49.61$, $SD=5.70$, $t(7)=12.32$, $p<.001$); PPA ($M=43.36$, $SD=5.40$, $t(7)=9.63$, $p<.001$); Parietal ($M=39.06$, $SD=7.84$, $t(7)=5.08$, $p<.01$).

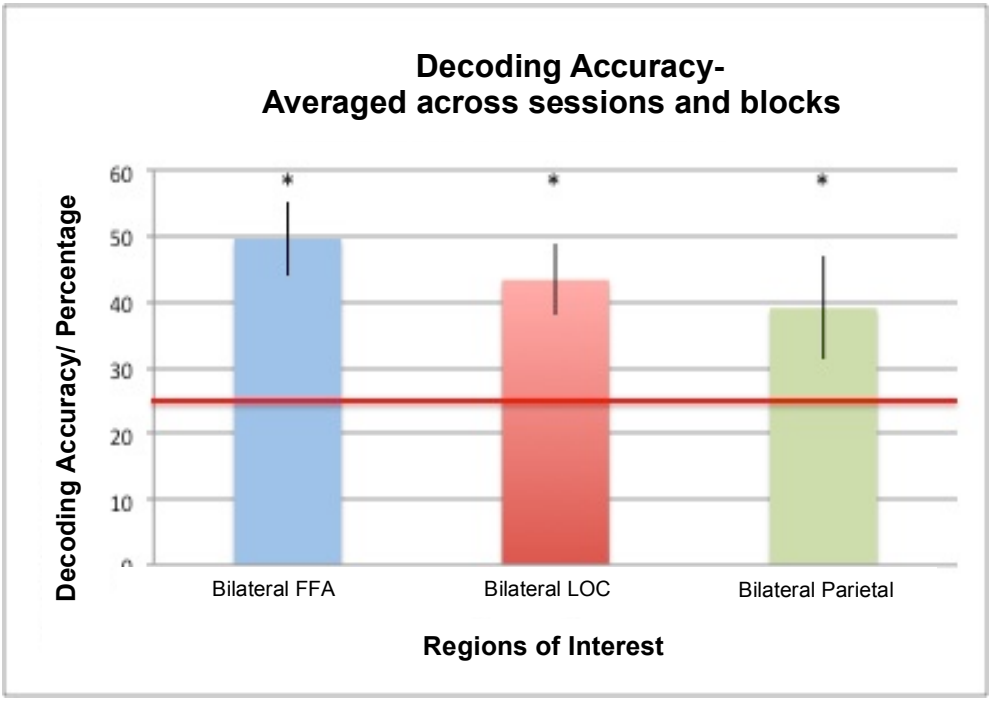

Fig 5-3. Participant-averaged decoding accuracy for the 3 ROIs averaged across sessions and blocks. Chance-level decoding is at 25% (horizontal red

line). The error bars show one standard error of the mean. Asterisks indicate decoding accuracy was significantly above chance. Error bars indicate ±1 SEM.

I next examined whether there was a change in decoding accuracy with time, over the sessions of the experiment. I performed a repeated measures ANOVA with the factors of session (1-4) and ROI (Bilateral FFA, Bilateral LOC, Bilateral Parietal). There was no main effect of session ($F(3,21)= 0.089$, $p= 0.97$). There was a main effect of ROI $F(3,14)= 6.29$, $p= 0.011$). There was no interaction between session and ROI $F(6,42)= 1.86$, $p= 0.11$). Thus, decoding accuracy remained the same over the course of the sessions in the experiment.

I further investigated a possible influence of time by looking at decoding accuracies over the 8 blocks *averaged across sessions and participants*. (see Fig. 5-3). A repeated measures ANOVA was performed, with the factors of ROI (Bilateral FFA, Bilateral LOC, Bilateral Parietal), and time (8 blocks, averaged across the 4 sessions). There was a significant main effect of ROI ($F(2,14)= 6.29$, $p=0.011$). There was no main effect of time ($F(7,8)= 2.14$, $p= 0.056$). This might suggest there was a decrease in decoding accuracies with time, although there was no significant interaction between ROI and time ($F(14, 98)=1.54$, $p=0.11$).

A paired t-test (2-tailed) comparing decoding accuracy over the 1st four blocks as compared to the 2nd 4 blocks revealed a significant decline in decoding accuracy for bilateral LOC ($t= 3.16$, $p=0.016$) and bilateral parietal ROIs ($t=2.61$, $p=0.035$). There was no significant decline in decoding accuracy for Bilateral FFA ($t= 2.09$, $p=0.075$). (Figure 5-4).

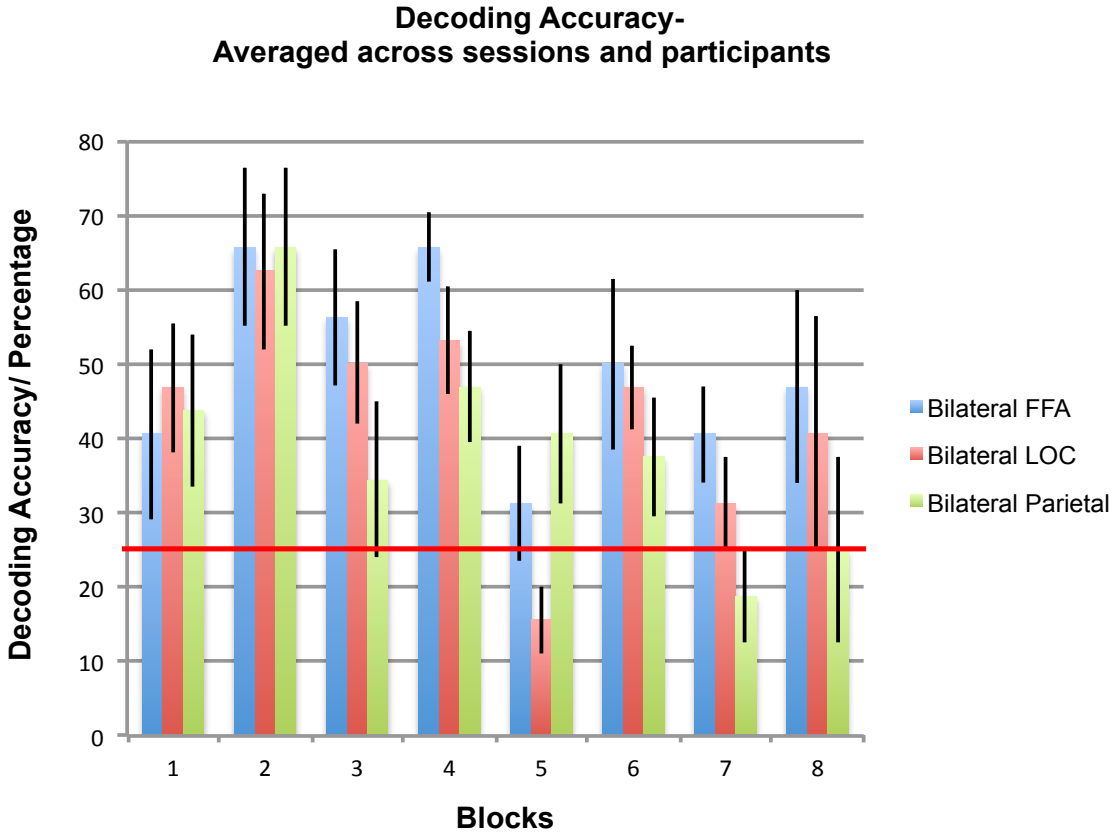

Fig 5-4. Participant-averaged decoding accuracy per block, averaged over the 4 sessions .Chance-level decoding is at 25% (horizontal red line). Error bars indicate ±1 SEM.

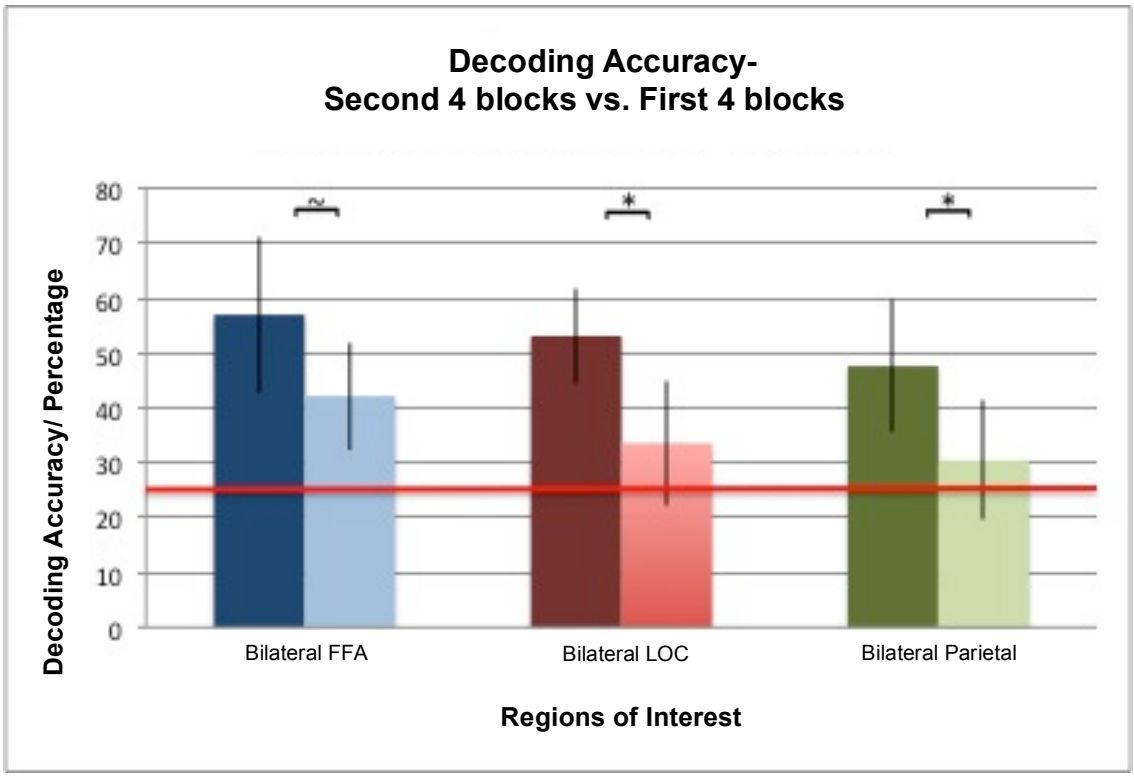

Fig 5-5. Decoding accuracies during each session, comparing the 1st four blocks with the 2nd four blocks. Chance is at 25% (horizontal red line). The columns in dark/solid colours represent decoding accuracy over the first four blocks, averaged across all sessions; the lighter/gradient-filled columns represent decoding accuracy over the second 4 blocks, averaged across all sessions. Decoding accuracy in bilateral LOC and Parietal ROIs was significantly higher during the first half of each session, as compared to the second half of each session. Error bars indicate ±1 SEM. Asterisks indicate significant differences in decoding accuracy, between the 1st four and 2nd four blocks.

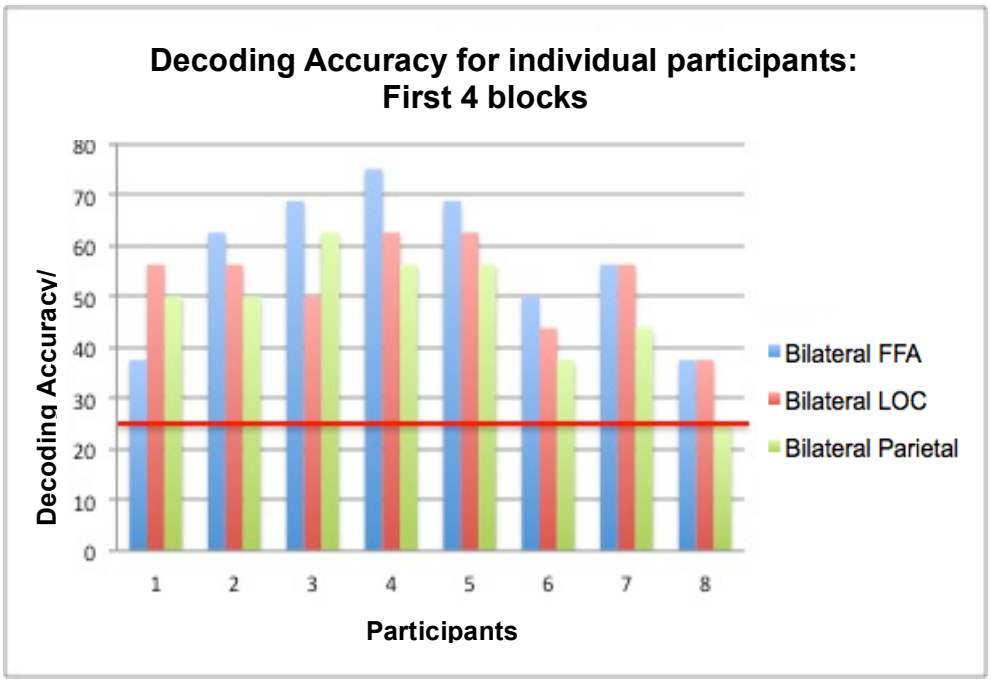

Figure 5-5A

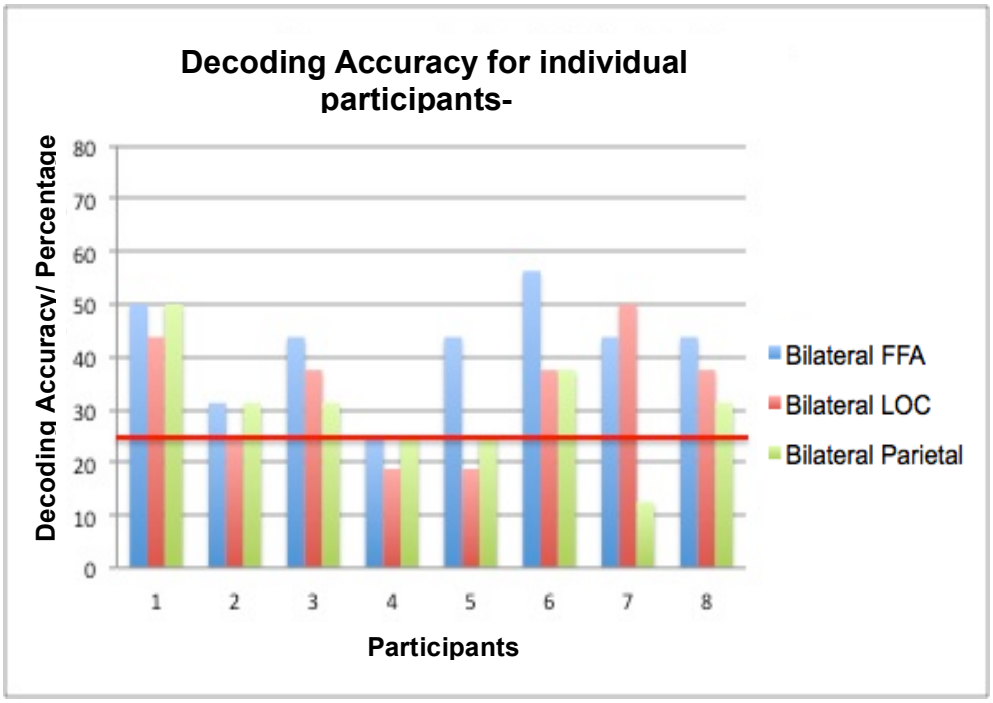

Figure 5-5B

Figs 5-6 A, B. Decoding accuracies for each participant, comparing the 1st four blocks (Fig. 5-6 A), with the 2nd four blocks (Fig. 5-6 B), averaged across all sessions. Chance is at 25% (horizontal red line). Note that this figure presents the same information as Fig 5-4, now dividing each session into halves.

3.3.2 N-back task reaction times

An assessment of reaction times on the n-back task was performed from the 6 participants on whom data was obtained, using an ANOVA, across sessions (1-4) and blocks (averaged over 1st four blocks, averaged over 2nd four blocks). This was to examine if fatigue may affect task performance across and within sessions (i.e. across blocks) (see Fig 5-7).

A main effect of blocks was observed ($F(1,5)=7.751$, $p=0.04$), with an increase in reaction times over the blocks (Figure 5-6). There was no effect of sessions ($F(3,15)=1.00$, $p= 0.42$), nor was there an interaction of blocks with sessions ($F(3,15)=0.49$, $p=0.70$).

In order to confirm the main effect of blocks (i.e. a change in reaction times over time), the block order was randomised (using a computationally generated random sequence) and then divided into halves. A repeated measures ANOVA was performed again, but revealed no main effects of session ($F(1,5)=0.89$, $p= 0.48$), or blocks ($F(1,5)=0.89$, $p= 0.11$), nor an interaction ($F(1,5)=0.64$, $p= 0.61$).

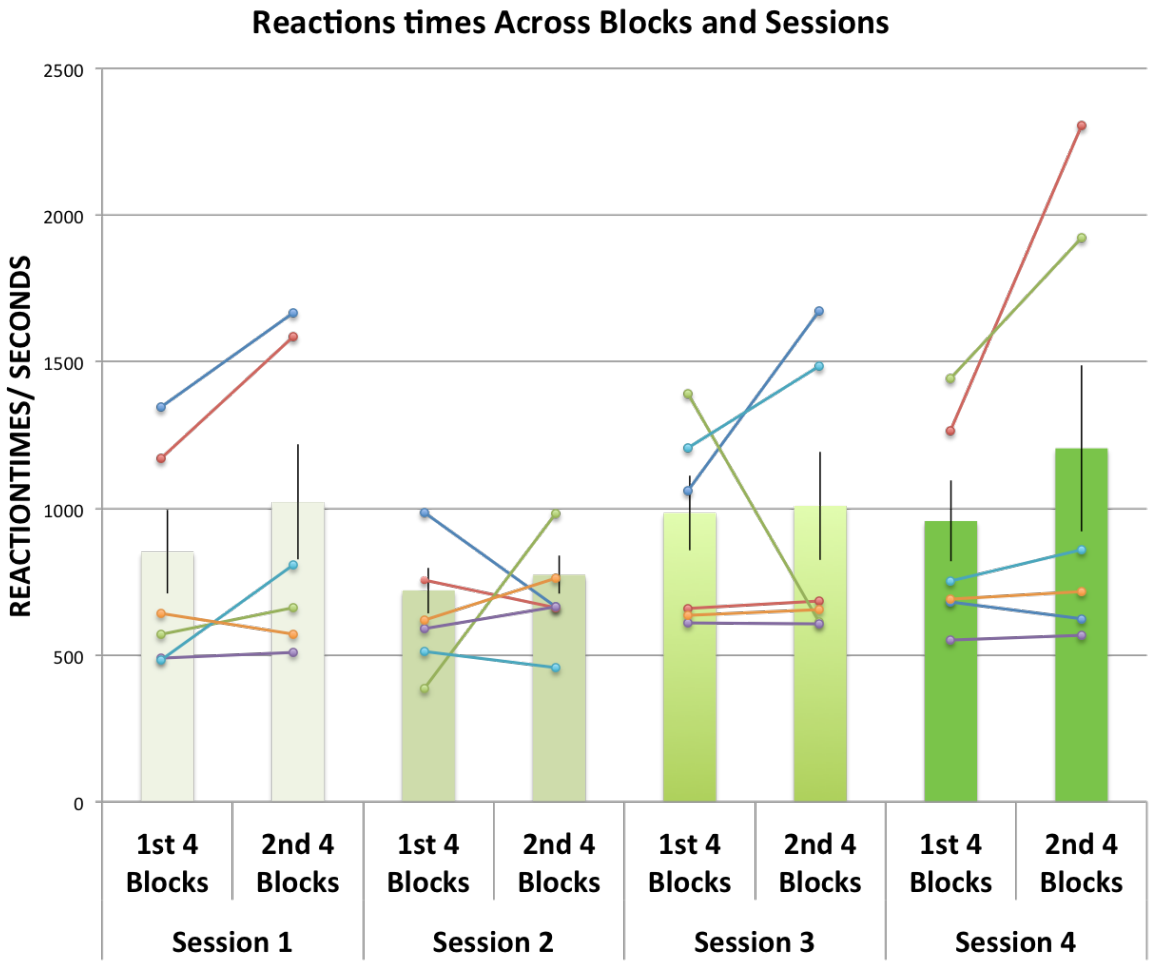

Fig 5-**7**. Graph showing reaction times averaged across participants for n-back task performance, for each session. **Sessions were divided further into the 1st four and 2nd four blocks to show the effects of experimental time on task performance. Error bars represent ± 1 SEM. Matched average reaction times for individual participants are shown for first four blocks and second four blocks of each session. A main effect of blocks was observed (F(1,5)=7.751, p=0.039), with an increase in reaction times over the blocks.**

5.5 Discussion

The purpose of this study was to develop a proof-of-principle for a rt-fMRI based BCI utilising online decoding of brain activations linked to the voluntary allocation of spatial attention. I report a novel online decoding method using rt-fMRI to identify the direction of covert spatial attention, using BOLD activation in higher order visual areas that have a role in the processing of spatial, and object-based information.

In this study I found online classification or 'decoding' of the voluntary deployment of covert attention to spatial location, could be implemented with a rt-fMRI set-up. A novel algorithm was implemented in MATLAB, using a 'winner-take all' decision rule, based on spatial quadrant-specific parameter estimates, and optimised using m-sequences to orthogonalise quadrant-specific stimulus presentation. Classification accuracies in the selected higher order visual ROIs i.e. FFA, PPA, parietal cortex, were significantly above chance in all 3 ROIs (all p's < 0.001); individual decoding accuracies were as high as 70%, during the first half of each experimental session. Using an assessment of participant reaction time on an interposed n-back task, the increased reaction times over the blocks within a session, were interpreted as suggesting participant fatigue may have contributed to a reduction of decoding accuracy towards the end of the BCI attention experimental session.

Attention is a versatile neural process, which is useful for the purpose of BCI development (15). In particular spatial attention is complex, engaging a number of cortical regions and combing both bottom-up and top-down processes (116,452–454). The outside world is spatially represented by internally maintained retinotopic maps, which have been demonstrated throughout the visual hierarchy. In addition to primary retinotopic cortex, retinotopy has been demonstrated in the dorsal (IPS (41)) and ventral processing stream (e.g. LOC) (48). In relation to higher order visual cortex, spatial selectivity has been documented in regions traditionally associated with feature and object based attention, including Occipital Face Area (OFA), FFA, Fusiform Body Area (FBA), Extrastriate Body Area (EBA), LOC, and posterior fusiform area (98). Object category and position may be jointly coded for in high order visual cortex such as the LOC (44,'45). Retinotopy has been demonstrated in

LOC (112), a region which also may serve as a point of integration between the ventral and dorsal visual processing streams (457). Golomb and Kanwisher (49) confirmed that topographical information relating to object position, is linked to retinotopic maps throughout high order object-selective regions, with top-down modulation of these neuronal populations resulting in spatiotopic representations. Further, attention may act to co-ordinate multiple perceptual streams related to separate visual and spatial properties of a stimulus, serving to integrate behaviourally relevant input (458). This body of work provided the evidence for the selection of the category and object-selective cortex such as FFA and LOC in this experiment.

The inclusion of temporal information by using m-sequences for stimulus presentation in our study served to disambiguate quadrant-specific brain activations in each of the higher order brain regions, thereby increasing classification accuracies during the selective deployment of attention to one of the 4 quadrants. The use of alterations in stimulus frequency has been previously examined in EEG and MEG, predicated on the improved temporal resolution of these modalities. The investigation of temporally selective attention has made a relatively small contribution to the existing literature. It is widely debated with regards to its exact neural mechanism e.g. is it sub-served by a discrete cortical locus, or does it function within existing components of attention, at a specific level of the attentional hierarchy. Coull and Nobre (459) identified the FEF and IPS, as being activated in fMRI and PET attentional studies involving manipulations of the time interval between stimulus presentations. Activations were greater when participants directed their attention across both space and time together, as compared to either dimension alone. Left hemispheric dominance was supported by collision-judgment tasks, with activations being identified in the left inferior- parietal cortex (supramarginal gyrus) (460,461). Cognitive manipulations of time perception on the other hand, activate the right hemisphere, with the right DLPFC and right parietal cortex being implicated (462–465). These findings suggest the potential overlap of discrete but integrated attentional pathways. Therefore, in addition to adding a temporal profile to the brain activations, the incorporation of a m-sequence into the stimulus presentation may have served to strengthen the neural activation related to the deployment of spatial

attention, as well as helping to segregate brain activations linked to the concurrent, but spatially separated presentation of visual stimuli.

I report four quadrant classification accuracies almost twice chance (44-50% accuracy in LOC/FFA, chance 25%). Previous studies using online non-invasive BCIs document accuracies of 41% for four quadrant decoding using MEG (420). Bahramisharief et al. (42) measured MEG occipital alpha band to examine covert attention over 16 polar eccentricities, in 11 participants. An offline analysis enabled the attended angle to be predicted in 7 out of the 11 participants. Kelly et al. (466) have also used occipital alpha band as the basis of a spatial attention EEG-BCI, reporting 73% classification accuracy over 10 participants, decoding left versus right covert spatial attention in offline analyses. Invasive BCI decoding using ECoG in non-human primates have reported accuracies of 67% (15). In relation to the use of online fMRI decoding of spatial attention, currently only one group have published using this method (10,44,311). Accuracy for four quadrant decoding was reported as being on average 78%, exceeding Kubler et al.'s previous suggestion of an operational accuracy of 70% for BCIs used for communication and control(78,467–469). Decoding was performed on neural activation in a priori determined retinotopic ROIs, each of which was quadrant-specific (43). Concurrent real-time feedback was also used to strengthen brain activation in the ROI corresponding to a cued direction of attention. With regards to online decoding in object based attention, Niazi et al. (470) used an online whole brain classifier to distinguish when participants were attending to a face or house stimulus, reporting accuracies of 77.6%. A recent study using object-based attention to drive a closed-loop neurofeedback BCI provided further evidence of using higher order visual regions such as FFA and PPA with rt-fMRI (323).

I extend the decoding approach, demonstrating for the first time proof-of-principle for a non-invasive cognitive BCI with rt-fMRI, utilising higher order visual brain regions, manipulating cognitive processes such as the covert deployment of attention. This BCI could serve as the basis for a communication interface in conditions in which patients have preserved higher order neural function, but are unable to speak or move to due to pathology affecting the brainstem or spinal cord. The methodological advance demonstrated in this study is important as it progresses BCI technology

from utilising primary sensorimotor cortex for motor rehabilitation to targeting internally driven cognitive processes. By accessing these more elusive neural activities, a attention-driven or cognitive BCI (9,15) also opens the door to cognitive enhancement and treatment of cognitive disorders. Further, I used stimuli that may be directly relevant in terms of a BCI device for practical use in a patient setting. Examples of using an attention driven BCI could include 'indicating' via images on a visual display, a particular body part that urgently needs medical attention, 'requesting' a food item or glass of water, or 'asking' for a particular individual. Each quadrant could provide a specific stream of information, which the individual could then direct their attention to as required. The potential benefit of targeting a focal brain region with a non-invasive BCI interface could also relate to the eventual translation into an implantable BCI. This might improve the signal-to-noise ratio by using intracranial recordings targeting neuronal populations identified and enhanced by prior use with a non-invasive BCI. Further such a device would have the important advantage of being portable and/or 'invisible'. Andersson et al. have previously reported on the use of rt-fMRI and cortical decoding as a prelude to classification using intracranial recordings with EcOG (10).

The accuracy with which the direction of attention to one of the four quadrants could be determined was as high as 60-70% during the first half of the scanning sessions, with a significant reduction in decoding accuracy during the second half of the scanning sessions. Sessions requiring the participants to deploy covert attention, or respond to a cue, made no difference to the accuracies. Interestingly, participants (e.g. participant 5) that exhibited effective (above chance) classification of their spatial attention deployment across the three brain regions examined in the first half of each scanning session (see Fig 5.6 A), performed less well during the second half of the experiment (see Fig 5.6 B), leading to an overall poor classification result across the three ROIs (see Fig 5.4). Examining the overall within-session classification alone (Fig 5-4) does not enable this difference to be detected. Possible reasons for this decline in decoding accuracies within may be related to fatigue, with there being a known reduction in goal directed performance e.g. top-down driven processes (459,460). Reaction times were examined as a surrogate marker of fatigue, with an increase in reaction times during the performance of the n-back task being noted during the course of a session, which was statistically significant (i.e.

F(1,5)=7.751, p=0.039). Mental fatigue, with subsequent impairment in complex task performance has been associated with a reduction in BOLD activation in underlying brain regions (461). Further, mentally fatigued individuals have been shown to have problems blocking out distracting stimuli (474). This consideration is important in the context of optimising a BCI imaging paradigm, in order to achieve the recommended criterion decoding accuracy of 70% (467–469). The implication from these findings is that an attention-driven BCI using realtime fMRI technology would be optimised by ensuring that sessions were kept short, for example 3-4 minutes (i.e. half of the session length used in this study). It is possible to envisage the long term utility of an attention-driven BCI for the selection of specific objects, using the allocation of spatial attention to navigate a visual-interface analogous to the BCI platform used here. The user would be able to interact rapidly with the interface to express a specific communication within the suggested time frame. In the event that the user-communication interface interaction was unsuccessful, or unable to be completed within a 3-4 minute 'user episode', there would be a requirement for a short period of rest before attempting usage again e.g. a few minutes. This would be a reasonable practicable approach to the use of such a BCI, except perhaps in situations potentially involving a patient's request for urgent attention, although this would be normally addressed through more standard patient care and support from carers/ healthcare providers.

Eye-tracking was not included in the fMRI part of the experiment, due to the complexity of the rt-fMRI set-up. Eye movements do represent a potential confound - brain activations associated with eye movements occurring in the cortical oculomotor network can overlap with those produced by shifts of spatial attention (475). In the EEG and ECoG field, it has been demonstrated that eye movements disturb decoding of attention-related brain activation, reducing classification accuracy to below chance level (476,477). I did use eye tracking in the offline version of this experiment and obtained similar classification accuracies in the same brain regions, whilst confirming an absence of excessive eye movements (see results Experiment 2).

5.6 Conclusions

One of the stated goals of BCI research is to provide a communication platform with real world applicability for patients who may lack the ability to move or speak e.g. stroke patients, or communicate with eye movements e.g. locked-in patients. In this study, I address this requirement by using an internal, higher order process which is not dependent on activation of motor or primary sensory cortex. Decoding the direction of spatial attention using combined univariate approaches on a single block basis was achieved significantly better than chance from higher visual areas including bilateral parietal cortex, LOC and FFA. Decoding accuracies approaching the recommended target of 75% were achieved during the first half of each scanning session. Decoding spatial attention, including temporal information and stimuli with real-world saliency may pave the way towards operational applications such as 'attention-based' BCIs for locked-in patients.

Rationale for Experiment 4

In the preceding three experiments, univariate and multivariate classification techniques were successfully applied to brain activations in higher order visual areas produced by cognitive tasks, culminating in an online implementation using an rt-fMRI experimental pathway. This work also provides proof-of-principle for a non-invasive attention driven BCI.

In the next experiment, I extend the established use of rt-fMRI with higher order visual areas, to examine the second BCI principle of a neurofeedback loop. By training participants to modulate brain activation in category-specific regions using neurofeedback, I investigate the potential for producing changes in visual perception, and linked changes in the functioning of the trained brain regions.

Abstract: Experiment 4

Functional plasticity in high order visual areas following neurofeedback training

Introduction

In the preceding experiments I explored and confirmed the ability to accurately decode brain activation in object specific brain regions such as LOC and FFA, in real time. Accessing the level of activation in these brain region in real-time may allow us to go further, beyond simply reading it out to manipulating the level of activity (see Fig 1-1, Chapter 1). This in turn could lead to changes in the functional, and structural plasticity of the brain region being modulated, potentially causing changes to linked behaviour. Neurofeedback-based training using rt-fMRI can be implemented to enable such manipulation, and when performed in higher order visual regions involved in complex visual processing, may lead to changes in perception.

I selected two discrete brain regions which have specific roles in object processing (i.e. FFA in face processing, PPA in house and place processing), as compared to the LOC which has a more general, overarching role in object processing. I hypothesised that learning to modulate the level of activation in two higher order visual areas, the FFA and PPA, using neurofeedback would change the dynamics of binocular rivalry for specific stimuli that activate these regions. I used a binocular rivalry (BR) paradigm where a face stimulus and house stimulus were presented simultaneously to each eye. I hypothesised that rivalry between the ensuing face and house percepts would be altered by learning-induced changes in the trained regions. To explore the mechanisms underlying neurofeedback learning, which involve top-down processes, I also examined functional, effective connectivity and structural changes in the SPL, a region previously implicated in neurofeedback learning.

Methods

Brain activation in FFA and PPA, and BR dynamics (cumulative dominance duration, cumulative switch frequency) were measured before and after neurofeedback training in ten participants. Cumulative dominance durations are a measure of the total time a face or house stimulus was perceived (duration of the percept). Cumulative switch frequencies measure the total number of times a face or house stimulus was perceived (frequency of percept). Participants were divided into two training groups who each learned up-regulation of the difference in signal between FFA and PPA (or vice-versa). Structural analysis in the trained regions and SPL were performed using tensor-based morphometry (TBM). Effective connectivity between SPL and the trained regions were examined using Dynamic Causal Modelling (DCM). A Canonical Variate Analysis (CVA) was performed to examine associations between the behaviour measures, and the structural and functional changes, before and after neurofeedback training.

Results

All participants learned to regulate the differential signal ($F(1,8) = 5.75$, $p = 0.04$). BR dynamics were altered in both groups with a change in the dominance durations of the percepts that were dependent on the signal they had trained on. Specifically, there was a *decrease* in the dominance duration of the house percept for the FFA minus PPA group, and of the face percept for the PPA minus FFA group ($t(9)=2.88$, $p=0.02$). Participants were also able to execute up-regulation during BR, causing a further shift in BR dynamics in a similar manner ($t(9)=4.76$, $p=0.001$). The difference in signal between FFA and PPA (or vice-versa) increased during concurrent up-regulation ($t(9)= 2.21$, $p< 0.03$, one-tailed t-test). A DCM analysis showed the training effect to be in the modelled bottom-up connections from ventral visual areas used to produce the training signal to SPL, and in SPL's self-connections. The CVA analysis confirmed an significant association between changes in dominance duration following neurofeedback training, and the structural changes in FFA, PPA and SPL.

Conclusion

Voluntary modulation of higher order visual areas can be learned using rt-fMRI neurofeedback, and leading to changes in visual perception. The effects of neurofeedback training of ventral visual areas were underpinned by functional changes in the trained brain regions, as well as changes in effective connectivity between a putative control region associated with visual perception, SPL, and the trained regions, FFA and PPA. Finally, individual differences in perceptual measures of BR were associated with structural changes in all 3 regions.

Key words

rt-fMRI, neurofeedback, differential signal, higher order visual areas, PPA, FFA, SPL, top-down control, DCM, perceptual biasing, CVA

6. EXPERIMENT 4

Functional plasticity in high order visual areas following neurofeedback training

6.1 Introduction

At any point in time, we are aware of only some of the information available to our senses. Our extraordinary ability to make sense of the visual scene is guided by the interaction between previously determined behavioural goals, as well as the salience of stimuli present in the scene (478). Familiarity and previous perceptual experience (479) facilitate perceptual processing, making it more rapid and efficient. Prior determination of internally generated task goals may lead to unconscious covert orienting in relation to basic stimulus features (480), as well as more abstract categorical information (481), invoking the use of working memory representations, and search-related attentional capture.

The neural basis for the selection of visual stimuli prior to processing has been linked with preparatory activity in category-specific visual brain regions, particularly when participants were provided with abstract strategies, prior to visual search for stimuli not previously seen before. Importantly, activation in primary visual cortex was not implicated, despite an established role of this region in visual imagery. A top down role for category-specific regions was therefore suggested, with internally generated neural activity in these regions biasing visual processing in favour of linked object categories (482). An key finding in this earlier work, was that preparatory activity was most effective in influencing processing when implemented at the level of category-specific brain regions

An unanswered question in the field is whether we can deliberately alter the likelihood of a stimulus entering awareness, by the direct approach of training brain areas involved in visual processing and awareness, such as category-specific regions. If visual awareness could be biased in this way, the likelihood of specific sensory interpretations occurring may be increased, with broader implications for manipulating visual perception. The visual system provides a useful platform to examine the manner in which sensory processes, which generate perceptual experiences, may be manipulated. Category-specific brain regions, attributed with processing of specific visual stimuli e.g. FFA and faces, PPA and houses, might be manipulated directly with neurofeedback using rt-fMRI. Altering the levels of brain activation in these regions using neurofeedback, can be explored in relation to visual processing, allowing use to investigate the potential for the unconscious biasing of visual perception.

Neurofeedback training with rt-fMRI, employs voluntary control of a visual interface, (e.g. fluctuating thermometer bar), presented to the participant while in the scanner. This is linked to the activation of target brain regions, by the participant using mental imagery (85), while being scanned. The neural substrates of mental imagery are thought to share considerable overlap with those engaged in the perception of the outside world (151,171) and behavioural execution (483,484), enabled by task-specific neuronal representations. Although mental imagery has been shown to alter perception in the short term (485), evidence of its effects following long-term practice are not clear. For example, successive training of mental imagery does not increase the vividness of the imagery or affect binocular rivalry (BR) between stimuli (486). In comparison, training of motor imagery produces task improvement and has been associated with functional and structural changes throughout the motor hierarchy (13,14,15). This suggest that visual perception may complex process with regards to its relationship with mental imagery.

The use of online data processing tools allows simultaneous measurement and feedback of brain activation to the participant while engaging in a neurofeedback task. The ‘training’ component of neurofeedback training requires voluntary control of the display, e.g. learning to consistently ‘push’ the level of the thermometer bar upto, or above a target level. The effect of training of primary visual areas in relation to

visual processes has recently been shown - Scharnowski et al. demonstrated improved visual detection during voluntary up-regulation of V1 activation following neurofeedback training (16), and Shibata et al. showed increased orientation discrimination following training of retinotopic cortex (areas V1/V2) (17). A further question is whether or not high order visual processing such as perception can be perturbed by engaging category specific visual areas in neurofeedback training. In this thesis I chose to examine the effects of neurofeedback training of higher order visual areas on BR, a well evidenced perceptual task (490–493).

BR describes the phenomenon where incompatible stimuli are presented separately and simultaneously to each eye of a participant (494,490). The stimuli presented for rivalry may be as simple as orientation gratings, which are marginally different from each other, or may be more complex stimuli such as natural scenes presented to one eye, and faces presented to the other. Regardless, the participant is consciously aware of only one of the perceptual stimuli at a time- the other stimulus is suppressed. This perceptual fluctuation is spontaneous and stochastic (495). Both top-down (i.e. higher cognitive modulatory) and bottom-up (i.e. salience-based) processes have been shown to be involved (496–498). Category specific brain regions, involved in perception of stimuli engaged in rivalry (e.g. FFA and PPA for face/ house rivalry) demonstrate changes in their levels of activation prior to and during phenomenological changes in BR (150,171), indicating the relationship of specific fluctuations in brain function to changes in perception.

An important consideration in the design of the study presented in this chapter was the role of top-down control in visual processing. In the context of perception, this top-down processing likely reflects feedforward and feedback interactions between multiple cortical regions. Activity in the frontoparietal network is likely to represent the top of the command hierarchy, with specific examples documented in both neurofeedback training of visual regions and in binocular rivalry. Scharnowski et al. (81,16) recently demonstrated increased functional and effective connectivity between a specific region of the SPL and retinotopic visual cortex following neurofeedback training of V1. Neurofeedback training is likely to engage a number of higher order processes, including those related to the neuroprosthetic control, agency, introspection and attention. Interestingly, regions of SPL has been shown to

be causally involved in BR. Kanai and Carmel (499,500), using transcranial magnetic stimulation, illustrated increased perceptual switching following stimulation of an anterior SPL site (aSPL), and converse following stimulation of a posterior SPL site (pSPL). The findings were interpreted within a Bayesian framework, with aSPL setting up perceptual predictions, and pSPL establishing the related prediction error between the prior perceptual expectation and the 'posterior' sensory evidence received.

I therefore set out to firstly establish the effect of neurofeedback training of high order visual areas on visual perception, using a binocular rivalry paradigm, and secondly, to delineate possible mechanisms involved in top-down control of neurofeedback training. I used rt-fMRI with neurofeedback to train a group of healthy participants to differentially modulate the level of activation in visual category-specific brain regions (FFA and PPA). Changes in perception during in BR, were recorded, in relation to the perception of pairs of face and house stimuli.

I hypothesised that neurofeedback would enhance the effectiveness of imagery-based training on visual perception by enabling voluntary modulation of the level of activity in category-specific brain areas. I predicted that the effect of neurofeedback training would be to alter the baseline functional state of the trained regions, as well as underlying structure and functional connectivity.

Specifically in relation to functional connectivity I investigated the involvement of top-down modulatory regions implicated in neurofeedback and BR – I examined changes in effective connectivity between SPL regions and FFA and PPA using dynamic causal modelling (DCM), following neurofeedback training.

6.2 Methods

6.2.1 Participants

Ten neurologically normal adult volunteers (24–35 years of age; mean age = 28 years; 8 females) with normal or corrected-to-normal visual acuity participated in the experiment. Each participant gave written informed consent; the study was approved by the local ethics committee (UCL Ethics Committee code: 09/H0716/14) and passed a MRI safety medical screening approved by the Wellcome Trust Centre for Neuroimaging (where MRI scanning was performed).

6.2.2 Stimuli and Materials

All visual stimuli described below were generated and displayed via scripts in MATLAB created with Cogent extensions. Stimuli were presented on a viewing screen with visual angle 23 degrees by 17 degrees (30 x 26 using a LCD projector (LT158; NEC). The stimuli were 3 degrees (in a tiled frame of 4 degrees). They were visible to the participants in the scanner via a mirror-mounted screen on the top of the MR scanner bore, behind the head, at an optical distance of 52 cm. Face or house greyscale stimuli were presented, with 20 different exemplars in each of the two categories. Responses were obtain via a pair of custom-built, MR- compatible, response boxes.

6.2.3 Experimental Outline

The experiment was divided into multiple sessions with each participant being required to attend 6 consecutive scanning days. The participants were split into two groups, with 5 participants in the ‘face’ group and 5 participants in the ‘house’ training group. On Day 1, all participants underwent a pre-training high-resolution

structural scan, a functional localiser scan and pre-training BR sessions. On Days 2-4, participants underwent neurofeedback training sessions, for a total of 3 separate training sessions. On Day 5 a transfer session was performed. On Day 6, post-training BR was performed, as well as a post-training high-resolution structural scan.

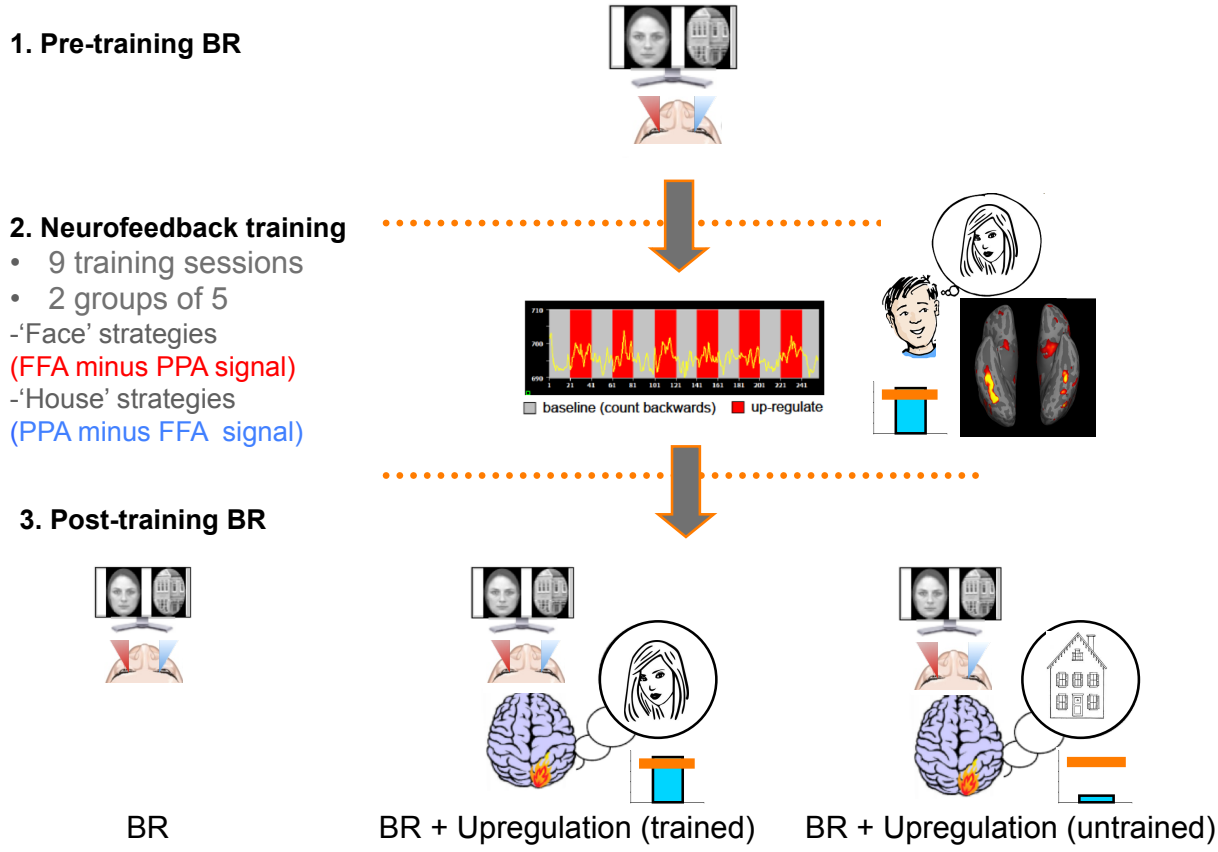

Fig 6-1. Experiment 4 experimental procedure schematic. The three stages of the experiment are outlined below:

Stage 1: *Pre-training BR*

Three sessions of BR were performed

Stage 2: *Neurofeedback training*

The 10 participants were separated into two groups prior to neurofeedback training. The 'face' group were explicitly instructed to use face-imagery strategies to 'drive up' a fluctuating thermometer bar (blue bar), up to a fixed mark (orange bar). The 'house' group were similarly instructed, but they were

to use 'house-imagery' strategies. The blue thermometer bar represented the difference in activation levels between the FFA and PPA ROIs – the face group trained on a FFA minus PPA signal, and the house group trained on a PPA minus FFA signal. Following the neurofeedback training sessions, the participants performed a transfer session during which they differentially modulated FFA and PPA as trained, but now in the absence of neurofeedback signal.

Stage 3: *Post-training BR*

This was separated into three sessions. The participants performed BR sessions (2), BR with concurrent up-regulation of target brain regions as trained (2), and BR with concurrent up-regulation of the target brain regions using non-trained strategies (2) i.e. the 'face' group used house-imagery related strategies (as illustrated), and the 'house' group used face-imagery related strategies.

6.2.4 FMRI scanning

All experiments were performed on a 3T Allegra head-only scanner, using a standard transmit–receive head coil. Functional data were acquired with a single-shot gradient echo planar imaging sequence (matrix size, 64_64; field of view, 192_192mm; isotropic resolution, 3 x 3 x 3 mm; 32 slices with ascending acquisition; slice thickness, 2 mm; slice gap, 1 mm; echo time (TE), 30 ms; repetition time (TR), 1920 ms; flip angle, 90°; receiver bandwidth, 3551 Hz/pixel). In the middle of each scanning session, double-echo fast, low-angle shot sequence (FLASH) field maps (TE1, 10 ms; TE2, 12.46 ms; resolution, 3 x 3 x 2 mm; slice gap, 1 mm) were acquired and used to correct geometric distortions in the images attributable to static field inhomogeneities.

6.2.5 Realtime fMRI set-up for neurofeedback

I used the Turbo Brain Voyager package (300) with custom real-time image export tools programmed in ICE VA25 (Siemens Healthcare) (301), and custom MATLAB-based scripts. Participants were shown visual representations of BOLD signal changes in brain regions that had been identified with a functional localiser scan i.e. target ROIs. Head motion was corrected in real time using Turbo Brain Voyager.

Real-time data preprocessing was performed in Turbo Brain Voyager and encompassed 3D motion correction with realignment to the preselected template, smoothing, incremental linear detrending of time series and statistical parametric mapping. The ROI time course(s) were extracted from the prescribed ROI masks, averaged and exported by the Turbo Brain Voyager. Signal drift, spikes and high frequency noise were further removed in real time from the exported time courses with the custom MATLAB scripts (302). The feedback signal was calculated and displayed to the participants in the form of a fluctuating thermometer bar, with a delay of 2 s from the acquisition of the image.

6.2.6 Binocular rivalry set-up

Inside the scanner, participants, wearing custom-made prism glasses, were shown two stimuli equidistant from a central viewing screen divider. To ensure fusion prior to the BR task blocks, two identical stimuli were presented to each eye; the images were moved with respect to the midline and each other until optimal fusion was achieved to produce a single image. During the viewing blocks a face stimulus and a house stimulus were presented in the left and right hemi-fields respectively - the stimuli were pseudorandomised with regards to which eye received the face or house stimuli; each viewing block was performed with a new pair of stimuli from the pool of 20 stimuli.

6.2.6 Day 1: Pre-training binocular rivalry

Both prior to and after neurofeedback training participants were tested on a BR paradigm. This involved a face image being presented to one eye, and a house image being presented to the other eye, simultaneously. Participants were instructed to indicate a perceptual shift only if the whole exemplar was perceived; any combination or 'patchwork' percept regardless of the predominance of the exemplar category was reported as a 'mixed' percept. The participants were instructed to switch as accurately and rapidly as possible between the three possible button presses linked to the three perceptual states (i.e. face percept, house percept, mixed percept). This resulted in measures of the cumulative frequency of switches to a specific percept, and the cumulative duration of the percept throughout the BR measurement period. Each measurement block lasted for 40s, followed by a rest period of 20s; a total of 4 blocks were performed per session, for a total of 3 sessions.

6.2.7 Day 1: Localiser

After the pre-training BR scan each participant underwent a functional localiser scan to identify FFA and PPA regions. The localiser session lasted for 12 minutes, consisting of 16 blocks of face stimuli, and 16 blocks of house stimuli, with 20 different exemplars per block. Each stimulus was shown for 600ms (400ms interstimulus interval). To ensure participants attended to the stimuli, they performed a one-back task in which they indicated with a button-press when a stimulus had been repeated (3 times/block).

Two contrasts were applied namely Houses vs. Faces, Faces vs. Houses (t-statistic threshold > 3.00); I then used the ROI selection plug-in for Brain Voyager QX, to demarcate and select the ROIs for each participant. The ROIs were defined along the ventral and lateral surfaces of the temporal lobe in proximity to the fusiform gyrus for FFA, and lateral to the collateral sulcus in the parahippocampal region for PPA respectively.

6.2.8 Day 2-4: Neurofeedback sessions

Each neurofeedback training session consisted of three neurofeedback scanning sessions; each session was made of 6 blocks lasting 60s, and consisted of 'up-regulate' periods (40s) followed by 'rest' periods (20s). During an up-regulation period, participants viewed a fluctuating horizontal red bar, and a fixed horizontal black bar, the latter being situated towards the top of the screen. Participants were advised that the fluctuating red bar was related to their brain activity; they were instructed to drive the red bar up to the level of the black bar by using mental imagery strategies (see below). Further, they were advised to try and maintain the red bar at that level, for as long as possible during the 'up-regulate' period.

Participants were also advised that there was a delay related to the training signal based on brain activity (i.e., the haemodynamic function of the BOLD signal) that they would be controlling, of approximately 6-8s. They were told not to be disappointed if there were moment-to-moment drops in the level of the fluctuating bar, but rather to focus on 'pushing' their brain activity in a consistent way in order to maintain a steady 'momentum' during the brain up-regulation blocks.

During the intervening rest periods, participants were instructed to perform a mental arithmetic task (serial subtraction of 7 from 100 i.e. 'serial 7s'). This orthogonal task ensured that brain activations during the 'rest period' (related mental arithmetic) were unrelated to those produced during the 'task period' (related face or house imagery).

6.2.9 Neurofeedback signal

The 'face' group were instructed to use mental imagery strategies involving face imagery, with possible examples of what might work (e.g. visualisation of familiar faces, famous faces). They were encouraged to identify a strategy that worked best for them, allowing a broad interpretation of the general face imagery instruction. Similarly, the 'house' group were instructed to use mental imagery strategies involving house imagery. They were also given broad examples of strategies, which

might be employed (e.g. visualisation of familiar buildings, the outside of their home), and encouraged to find strategies that worked best for them from session to session. Across both groups, it was open to the participants to utilise their own interpretation of the basic instruction.

The signal that participants were trained to regulate was derived from the ROIs identified during the Day 1 functional localiser. Specifically, during neurofeedback training, the fluctuating red bar was driven by brain activity in which the signal from PPA was subtracted from the FFA-derived signal (for the 'face' group) or the reverse subtraction (i.e., PPA – FFA) for the 'house' group). Participants were therefore trained to modulate a *differential* training signal, rather than the activity in a single region of the brain. The approach prevented participants from adopting a strategy that might modulate global BOLD activation rather than the desired ROI-specific changes.

There were therefore 5 potential activation states which could increase the difference between the two brain regions (R1 and R2), leading to up-regulation of the differential training signal i.e. R1 - R2. These could be (1) An increase in R1; (2) a decrease in R2; (3) a combination of the two; (4) a relatively greater increase in R1 as compared to R2; and (5) a relatively greater decrease in R2 as compared to R1.

6.2.10 Learning effect

The learning effect measures the change in BOLD activation in trained brain region/s across the neurofeedback training sessions. This involves calculating the mean percentage signal change (PSC) on each training run, and plotting it as a graph of mean percentage signal change against time. For my experiment, participants specifically learned to control the difference in activation between two brain regions. Therefore, the learning effect for each training run (3 per training day) for each participant was established by calculating the mean PSC based on the difference in brain activation between the two ROIs. Mean PSC was calculated in the following manner:-

PSC within ROI =

$$\frac{\textbf{Average Beta value in ROI during up} - \textbf{regulation Blocks}}{\textbf{Average Beta value for global constant during up} - \textbf{regulation bocks}}$$

x 100

Therefore:

PSC of training signal = (PSC ROI 1 - PSC ROI 2) per block

A group 'learning effect' was therefore represented by plotting for each of group, the average percentage signal change across the nine training sessions.

6.2.11 Structural scans for each neurofeedback session

A 4.5 min T1-weighted structural scan was performed at the beginning of each session (whole brain coverage; 3D FLASH; 1 mm isotropic resolution; matrix size, 256 x 224; field of view, 256 x 224 mm; 176 sagittal partitions; TE, 3.5 ms; TR, 9.5 ms; flip angle, 18°; readout bandwidth, 199 Hz/pixel; anatomical image data exported in near real time similar to Weiskopf et al. (301).

This anatomical image was used for co-registration of the in-session head position with the high-resolution T1-weighted structural MDEFT scan using TBV performed on the first scan day. From the resulting co-registration matrix, the position of the neurofeedback target ROIs in the current head position of the current run was determined. This ensured the same ROI was used in each of the sessions, on each of the scan days, for each participant.

6.2.12 Day 5: Transfer Session

On Day 5, following the final neurofeedback training session, participants were required to demonstrate self-regulation of differential brain activity, in the absence of a visual neurofeedback signal. This was in order to confirm that the participants had successfully learned to regulate brain activity.

There were 2 transfer sessions, each comprised of 6 blocks. Each block lasted 60s and consisted of an 'up-regulate' period (40s) followed by a 'rest' period (20s). During the up-regulation period, participants were presented with the word 'regulate' at the centre of the screen. Participants were required to try and drive their brain activity 'up', using mental imagery strategies as successfully used during the neurofeedback training sessions, but now in the absence of a feedback signal.

6.2.13 Day 6: Post-training binocular rivalry

Following the neurofeedback training sessions, all participants performed BR again, with the identical experimental set-up described above as for pre-training BR. This was to establish what the effects of neurofeedback training were on visual perception (i.e. the 'behavioural' effect). Three different BR conditions were performed in a counterbalanced order across all participants, consisting of 2 of the total 6 sessions. The 3 conditions were the following:

i) BR no up-regulation. The instruction here was identical to the pre-training BR, and will be referred to as 'baseline'.
ii) BR with concurrent brain modulation using the trained strategy**.** The 'face' group were instructed to use their individual face imagery strategies that worked best during the training and transfer sessions, while simultaneously engaging with the BR task. Similarly, the 'house' group were instructed to use their individual house imagery strategies, which worked best during the training and transfer sessions.

iii) BR with concurrent brain modulation using the non-trained strategy. This condition required the participants to generate imagery in a similar fashion to their training sessions, except with the explicit instruction that had been given to the *other* group i.e. the 'face' group were instructed to use house imagery strategies, and the 'house' group were instructed to use face imagery strategies. This condition enabled a comparison between the effect of trained up-regulation on BR, and up-regulation without prior training.

These 3 post-training BR sessions were pseudo-randomised and counter-balanced across the participants to address potential order effects and fatigue.

6.2.14 High resolution structural scans

A whole brain high-resolution T1-weighted structural scan was performed prior to training, and repeated once again following training. This was in addition to the 'short' structural scans (described above) performed on each neurofeedback scanning day. The scan was a 3D modified driven equilibrium Fourier transform (MDEFT) scan, 1mm isotropic resolution; matrix size, 256 **x** 240 mm; field of view, 256 **x** 240 mm; 176 sagittal partitions; TE, 2.4 ms; TR, 7.92 ms; inversion time, 910 ms; flip angle, 15°; readout bandwidth, 195 Hz/pixel; spin tagging in the neck with flip angle 160° to avoid flow artefacts for superposition of functional maps (501).

6.3 Analysis

6.3.1 Behavioural data – binocular rivalry

During the BR sessions, dominance durations for a particular percept, measured as the time taken from a switch to that percept, to another percept was collected from each participant. This also enabled switch frequencies to be calculated, based on

using switches *to* a specific percept only. Using custom- written MATLAB scripts, these measures were calculated per block, and session, to produce cumulative time measures for pre-training and the three post-training BR conditions. This yielded a cumulative dominance duration that was equal to the total amount of time the stimulus was perceived, averaged across blocks, and cumulative switch frequencies that was equal to the total number of switches to a particular percept.

Between condition changes in cumulative durations, and cumulative switch frequencies, for each percept, across groups were also calculated. Specifically, these were changes in the training-strategy related percept, the non-training strategy-related percept, and the mixed percept. For inferential statistical analyses of these dependent variables I used SPSS 21 (IBM Corp. Armonk, USA), to perform ANOVAs and follow-up planned t-tests, which were two tailed unless otherwise stated.

6.3.2 Brain imaging data – Binocular rivalry

Participant specific functional data were first analysed using SPM8. To allow for T1 equilibration, the first five images of each session were discarded. Preprocessing of the data involved mean bias correction, realignment (of each scan to the first scan of each run), unwarping, and co-registration of the functional data to the structural scan. The data were smoothed with a 6 mm FWHM Gaussian kernel and high-pass filtered (128s cut-off) to remove low-frequency noise including differences between sessions, while at the same time preserving as many of the spontaneous fMRI fluctuations as possible (433). Session specific grand mean scaling was applied, with no global normalisation, when preparing the GLM. Movement parameters from the three directions of motion and three degrees of rotation were included as confounds in the analysis of the imaging data.

6.3.3 Offline ROI analysis

6.3.3.1 Fusiform Face Area and Parahippocampal Place Area

Neurofeedback

BOLD changes across the 9 training sessions were modelled through the creation of a GLM, with regressors for each of the 9 sessions. Box-car functions were created for the four up-regulation blocks during each training session, convolved with the canonical HRF. Six regressors for movement and a global constant were also included. Rest was modelled implicitly. For each participant, beta values were extracted, averaged across all the voxels in the ROI masks (created using functional activations in FFA and PPA during the functional localiser). Mean PSC calculated, adjusted for the global brain signal (see equation above). For each participant, differential mean PSC between the two trained ROIs (i.e. training strategy related ROI and non-training strategy related ROI) was calculated across sessions. Using these values, the average mean PSC across participants, across the training sessions, and two subsequent transfer sessions were calculated.

Binocular rivalry

In order to determine BOLD signal changes during BR in the trained ROIs, FFA and PPA, boxcar functions were created modelling onset of the BR block, convolved with the canonical. HRF. This was done for each condition – pre-training BR (‘pre BR), post-training BR (‘post BR’), post-training BR with concurrent trained up-regulation (‘post BR + trained up-regulation’), post-training BR with concurrent non-trained up-regulation (‘post BR + non-trained up-regulation’). A GLM was performed at the single participant level. Beta values were extracted from each voxel for each of the trained ROIs, for each condition, then averaged and a mean PSC calculated, adjusted for the global brain signal (see equation above). Specific changes across groups in the training-strategy related ROI, and the non-training strategy-related ROI, between pre-training and post-training BR were examined. ANOVAs, and

follow-up planned t-tests (two tailed t-tests were used unless otherwise stated) were used to examine between condition changes in brain activation.

6.3.3.2 Superior Parietal Lobule

I examined SPL regions previously documented in the literature, in relation to neurofeedback learning and BR.

Neurofeedback

A pre-specified 4mm sphere (centred on MNI coordinates: x=22, y=-58, z=64 (81,16)) was used, based on previous work confirming activation during neurofeedback learning (81,16). In a an analogous manner to the method described above for FFA and PPA mean PSC for this ROI was extracted for each of the 9 training sessions, and the two transfer sessions, and plotted as an 'activation' curve (cf. 'learning' curve). A second level analysis was also performed as a sanity check to confirm significant activation in SPL during neurofeedback learning, using a small volume correction with the aforementioned co-ordinates.

Binocular rivalry

During BR, activation in pre-specified 4 mm spheres with central MNI co-ordinates of x= 36, y=-45, z=51 and x= 38, y=-64 z=32 was examined, based on previous work demonstrating activation during BR (499). At the second level, a small volume correction was applied across the ROIs examined for each of the BR conditions, as well as for the following contrasts: pre BR vs. Post BR, Post BR vs. Post BR + trained up-regulation, Post BR vs. Post BR + non-trained up-regulation.

6.3.4 Brain imaging – structural analysis

Two techniques were used for longitudinal structural analyses. In order to proceed with structural analysis, the structural brain data (collapsed across time) was first optimally aligned. For each participant, DARTEL was used, an algorithm for diffeomorphic image registration which optimises interparticipant registration (294).

For each participant, their preprocessed high-resolution T1 structural images were orientated to specifically place the anterior commissure, at the MNI origin. Segmentation of the images into grey matter, white matter and CSF in native participant space was then performed, using the tissue probability maps.

For each participant, a rigid body transformation based on six parameters was performed, with the DARTEL algorithm passing through six iterations of alternation between the creation of an average template of all the individual maps (i.e. a median image), and warping the individual participant images to the median image. The result of this was then used to create a series of flow fields per participant, which parameterised the deformation through non-linear registration of each individual image to the template for grey matter and white matter. The final average template was registered to the MNI 152 template brain.

TBM was then performed. TBM makes inferences based on Jacobian matrices of the deformation fields – these values correspond to the relative volumes of tissue before and after warping. These matrices vigorously track local structural changes such as contraction (Jacobian determinant <1) or expansion (Jacobian determinant >1). The use of Jacobian determinants makes TBM advantageous for small numbers of participants. For each participant, Jacobian determinants of the deformation fields were first calculated, and then the images were smoothed with a 6mm FWHM isotropic Gaussian kernel.

ROI masks were used, corresponding to the two trained ROIs (FFA and PPA), and aSPL. These ROIs were spheres centred on co-ordinates that demonstrated the highest functional activity across training (6mm for FFA, PPA; 4mm for aSPL). Values corresponding to changes in volume for pre-neurofeedback training vs. post-

neurofeedback training were extracted from these ROIs; ANOVAs, and follow-up t-tests (two tailed) were then performed to establish if specific brain regions changed significantly before versus after training.

6.3.5 Dynamic Causal Modelling

6.3.5.1 Overview

DCM is a framework for creating, estimating and comparing models of brain connectivity. DCM is 'causal' in the sense of control theory (502), and employs a forward model in which experimental manipulations cause changes in neural dynamics that may propagate to other neuronal populations. In turn, these dynamics are modelled as driving the haemodynamic response and BOLD signal. The basic procedure is to construct a simple neuronal model of interacting cortical regions, coupled with a realistic biophysical model of haemodynamic response, for each participant. Estimating the parameters of this model enables inferences to be made about the casual interactions between regions.

6.3.5.2 Introduction to DCM for fMRI

DCM for fMRI is a framework for inferring neuronal connectivity from regional haemodynamic BOLD timeseries (502,503). It rests on a two-part model. To begin with, change in the neuronal states of the brain regions of interest (Δz) with time (Δt) due to an experimental manipulation, is modelled with a differential equation:

$$\dot{z} = \frac{\Delta z}{\Delta t} = f(z, u, \theta^n)$$

Represented diagrammatically (annotations explained in text below diagram):

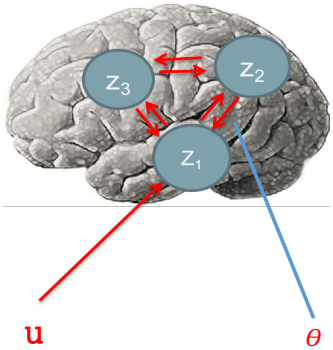

The neuronal model, f, predicts the rate of change in neuronal activity over time, $\dot{z}$, given external experimental inputs u, and connectivity parameters θ^n. These

connectivity parameters include the strength of connections or 'coupling' within and between regions. *z* represents the brain regions or 'nodes' being modelled. (NB: In the neurofeedback experiment presented the brain regions would be SPL, FFA, PPA, and the experimental input in relation to the neurofeedback training sessions would be the task of controlling the visual interface i.e. fluctuating thermometer).

$$y = g(z, \theta^h)$$

The observation model, g, is a detailed biophysical model which predicts the BOLD timeseries that would be observed, given the neuronal activity z and observation parameters θ^h. More specifically, the neuronal model has the following form:

$$\dot{z} = \left(A + \sum_{i=1}^{m} u_i B^i\right) z + Cu$$

The change in neuronal system with respect to time, $\dot{z}$, is a function of connectivity parameters A, B and C. Matrix A represents the strength of connections within and between regions in the absence of any experimental input These connections may be modulated by experimental input u, with matrix B^i representing the combination of the effect of experimental input u_i on each connection i.e. the change in intrinsic 'coupling' as a result of experimental input. Activity in the network is driven by experimental stimulation according to matrix C, which represents the influence of experimental input on each region. (NB:- In this study, this would be a visual input i.e. fluctuating thermometer effecting those brain regions directly engaged in the task of neurofeedback modulation, FFA and PPA). For each experimental condition $i = 1 \ldots m$, timeseries u_i represents the onsets of experimental stimulation. *z* represents the brain regions or 'nodes' being modelled. (NB: for the original theoretical and empirical work in DCM see (503–506)).

In this basic formulation, the level of activity in each ROI is represented by a single number or state, which forms the entries of vector z. Here, I used an extended two-state version of this model (504), in which each region was modelled as having an excitatory and inhibitory population of neurons, and long-range connections were

modelled as excitatory, in keeping with known neurobiology. With this extended model, the off-diagonal elements of matrix A represent long-range excitation, and the diagonal elements of matrix A represent self-inhibition within each region (i.e. the influence of each region's inhibitory population on its excitatory population).

As DCM is a Bayesian scheme, each parameter of the model was constrained by a prior distribution. For the neuronal part of the model, I used conservative 'shrinkage priors' which presume that coupling parameters are zero at the outset. For the haemodynamic part of the model, DCM is supplied with priors based on empirical measurements. The parameters of the forward neuronal model $\theta^n = \{A, B, C\}$ are combined the parameters from the observation model, and Bayesian inversion performed. This provides the evidence (i.e. marginal likelihood) of the model, as well as the posterior distribution of its parameters (i.e. neuronal coupling measures).

6.3.5.3 SPM Analyses

Prior to ROI creation, the functional images were smoothed with a 6mm Gaussian kernel to compensate for inter-participant variability. In preparation for the DCM analyses, the imaging timeseries from the 9 training sessions were concatenated. Onsets were determined for each of the 'up-regulation' blocks (4 per blocks per session, 3 sessions per training day), and a single regressor for training was created. Other regressors included those for session, and day. Six movement regressors were also included.

6.3.5.4 Regions of Interest selection and time series extraction

Six mm ROI spheres were created for both FFA and PPA for each participant, using the masks created following the functional localiser (contrasts of Face > House, and House > Face respectively). Timeseries from these volumes of interest (VOIs) were summarised using the SPM 12 eigenvariate function. I extracted each participant's principal eigenvariate around the participant-specific local maxima activation nearest to the peak voxel, within the mask. SPL VOI creation, and timeseries extraction were performed as a single step using the SPM 12 eigenvariate function, with extraction

around participant specific local maxima nearest to the voxel co-ordinates previously established (i.e. SPL; x= 22, y= -58, z=63). All voxels used to create the eigenvariates, were significant at $p < 0.05$ uncorrected and effects-of-no-interest such as motion and session means were regressed out. The single timeseries representing the activity in each region was then used to fit DCMs.

6.3.5.5 DCM – Specification

DCMs were created and estimated (503) using DCM 12 as implemented in SPM 12. All DCMs were deterministic, bilinear, 2-state models with mean centred inputs, and self-connections enabled. I created a model for each participant with three regions – a Training-strategy related ROI, a Non-training-strategy related ROI and SPL. All entries in matrix A were enabled (had non-zero prior variance), making all three regions fully connected. There were two experimental inputs (u) – Task and Training. Task, u_1, was a timeseries consisting of ones for each time step when the task was being performed and zeros for intervals between trials. Training, u_2, was a timeseries with zeros for sessions 1-3 and ones for sessions 6-9, modelling an increase in response due to training. All entries in matrix $B^{(1)}$ were fixed at zero, meaning there was no modulatory influence of Task on any connection. All entries in matrix $B^{(2)}$ were enabled, meaning that Training could modulate any connection. Matrix C was configured such that Task had a driving influence on the trained brain regions only i.e. not SPL.

6.3.5.6 DCM - Estimation

The BOLD timeseries were high-pass filtered, and motion confounds were removed. The parameters of each participant's model were estimated using an algorithm based on variational Bayes. This provides an estimated posterior distribution for each parameter, in addition to the free energy, which is a lower bound on the log model evidence. The free energy may be thought of as the accuracy of the model minus its complexity and enables model comparison.

After fitting each participant's model, post-hoc model reduction (506) was used to prune away any parameters which did not contribute to the free energy. In brief, in each iteration of the post-hoc search procedure, one or more parameters were disabled by setting their prior variance to zero. The free energy of this reduced model was rapidly calculated (using a generalisation of the Savage-Dickey ratio). If the reduced model did not have a worse free energy than the full model, it may be concluded that the disabled parameters were not contributing to the free energy. The algorithm repeats by disabling different combinations of parameters until an optimal is found.

The reduced models from each participant were averaged to achieve a group level reduced model ('Bayesian Parameter Average' or BPA). I then analysed the posterior estimates of the parameters in this reduced model, particularly to identify any effect of neurofeedback training.

6.4 Examining individual differences using Canonical Variate Analysis

In the final step, I sought to examine the effect of neurofeedback training on the combined measures of behaviour, functional BOLD changes during neurofeedback training, and structural changes in the trained regions, including the putative higher order control region, SPL. Statistical characterisation of the putative effects of behavioural changes associated with neurofeedback training (i.e. the change in BR dominance durations for all three percepts, comparing pre-training BR with post training BR with and without concurrent upregulation of trained brain regions), functional (i.e. BOLD activation in the FFA, PPA, and SPL ROIs during the neurofeedback training sessions) and structural (i.e. measures of volume change in structural imaging performed before and after training in FFA, PPA and SPL) was achieved through analysis of the individual subjects' measures using standard multivariate techniques (Canonical Variate Analysis (CVA), also known as multivariate analysis of variance, or ManCova .(507,508) CVA enables statistical inferences to be made about associations between the imaging data and behavioural data that are distributed over variables. It was chosen for analysis of this dataset

because it can accommodate statistical dependencies between multivariate predictor variables (behavioural changes) and multivariate outcome variables (functional or structural measures). This meant that neither the behavioural nor imaging data had to be examined in isolation, which had the profound advantage that distributed changes could be identified (while minimising the multiple comparison problem).

CVA was used to assess and characterise the presence of a mapping between the imaging dataset and the predictors of behavioural change. The objective of the CVA was to find the linear combination of outcome variables that was best predicted by a linear mixture (contrast) of structural or functional components. The weights of these linear combinations are called canonical vectors. The canonical variates of the outcome and predictor variables are the expression of each canonical vector in each subject. Other quantities generated by CVA include Bartlett's approximate chi-squared statistic for Wilks' Lambda and its associated significance, or p-value, which test for the significance of a linear mapping or correlation between the canonical variates. In other words, one or more pairs of canonical variates show a significant statistical dependency.

6.5 Results- Behavioural

6.4.1 Binocular rivalry

For each participant, cumulative dominance durations and cumulative switch frequencies were calculated based on the dominance durations and switch frequencies per block, performed during pre-training BR as well as for post-training BR. Three percepts were reported 'face', 'house' and 'mixed'. These were subsequently binned according to group membership in the following manner- the **'Trained' percept** (i.e. the response associated with the trained strategy/region – Face percept for the Face/FFA trained group, House Percept for the House/PPA trained group), the **'Non-trained percept'** (i.e. non-congruent with the trained strategy/region- House percept for the Face/FFA trained group, Face percept for the House/PPA trained group), and the **'Mixed percept'** (representing any combination

of the Face or House percepts). In post-training BR, in addition to performing baseline BR, participants also performed BR while simultaneously up-regulating brain activity as performed during neurofeedback training and transfer session. Additionally, non-trained up-regulation was attempted by all participants while performing BR.

6.4.2 Durations

Planned paired t-tests were used to examine cumulative dominance durations for each of the three percepts (grouped into training-strategy related percept, non-training strategy related percept, and mixed percept) specifically comparing the pre-training baseline state with the three post-training states, which were 'baseline' BR, 'BR with trained up-regulation', and 'BR with non-trained up-regulation'. (See Figs 6-2 and 6-3).

Comparing dominance durations from baseline BR with those from post-training BR, there was a significant reduction of the non- training strategy related percept ($t(9)= 2.88$, $p=0.018$), and a significant increase in the duration of the mixed percept ($t(9)= 2.74$, $p=0.023$). No significant change was evident for the training strategy related percept ($t(9)=0.46$, $p=0.66$). (See Fig. 6-3).

Comparing baseline BR with post-training BR + trained up-regulation, there was a very significant reduction in the cumulative dominance duration of the non- training strategy related percept ($t(9)=4,76$, $p=0.001$), and a significant increase in the duration of the mixed percept ($t(9)=2.68$, $p=0.03$). No significant change was produced for the training strategy related percept ($t(9)=0.53$, $p=0.61$). (See Fig. 6-3).

Comparing baseline BR with post-training BR + non-trained up-regulation there was a significant reduction in the duration of the training strategy related percept ($t(9)= 2.41$,$p=0.04$), and a significant increase in the duration of the mixed percept ($t(9)= 3.60$,$p= 0.006$). (See Fig. 6-3).

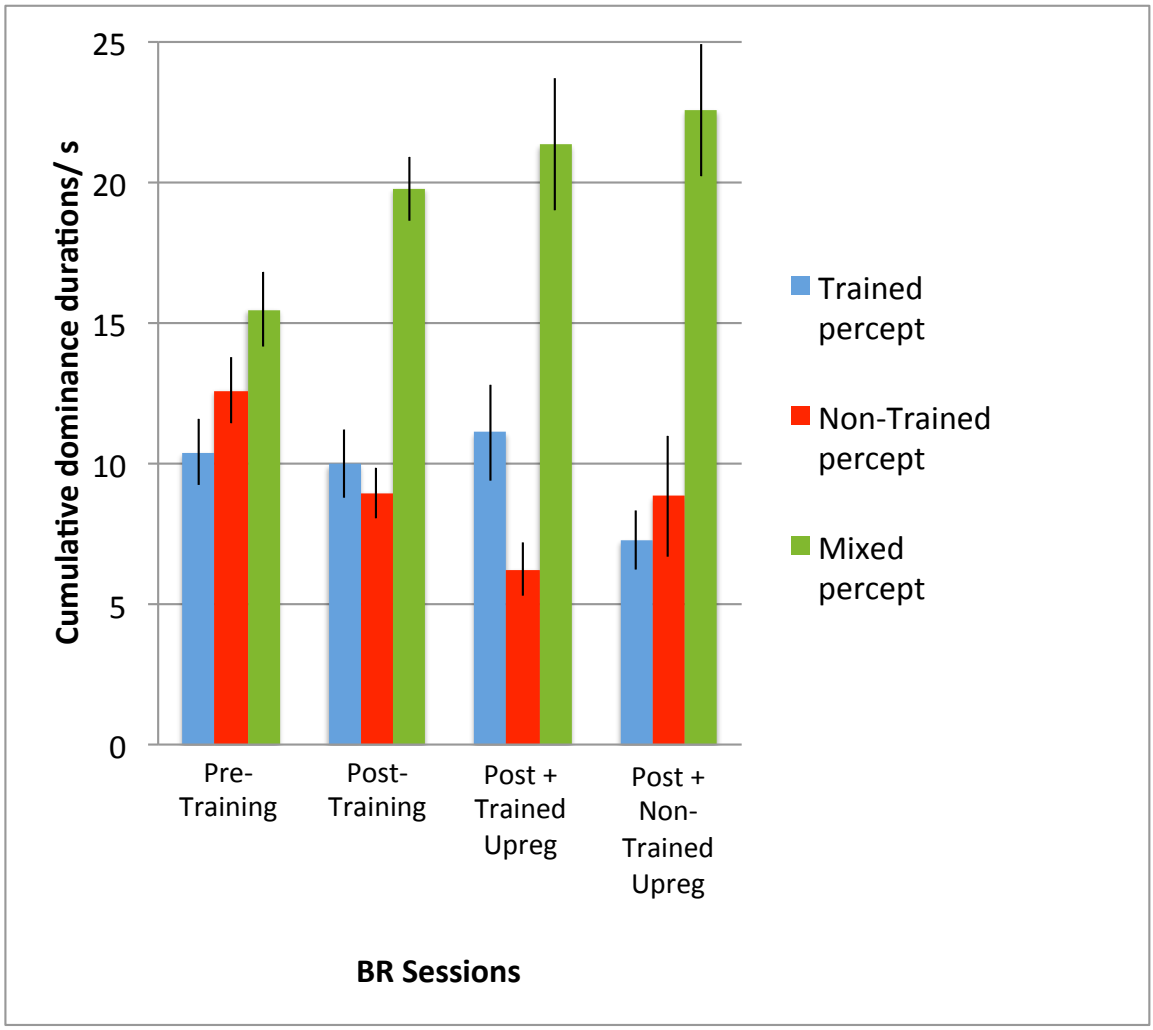

Fig 6-2. Binocular rivalry cumulative dominance durations across participants. Cumulative dominance durations (i.e. during a session, the total duration a percept linked to training strategy, not linked to training strategy or mixed combination of the two, was perceived) across participants for pre-training BR, and the three post-training BR sessions – Post training, Post training BR with concurrent trained up-regulation, and Post training BR with concurrent non-trained up-regulation. Error bars represent standard error of the mean.

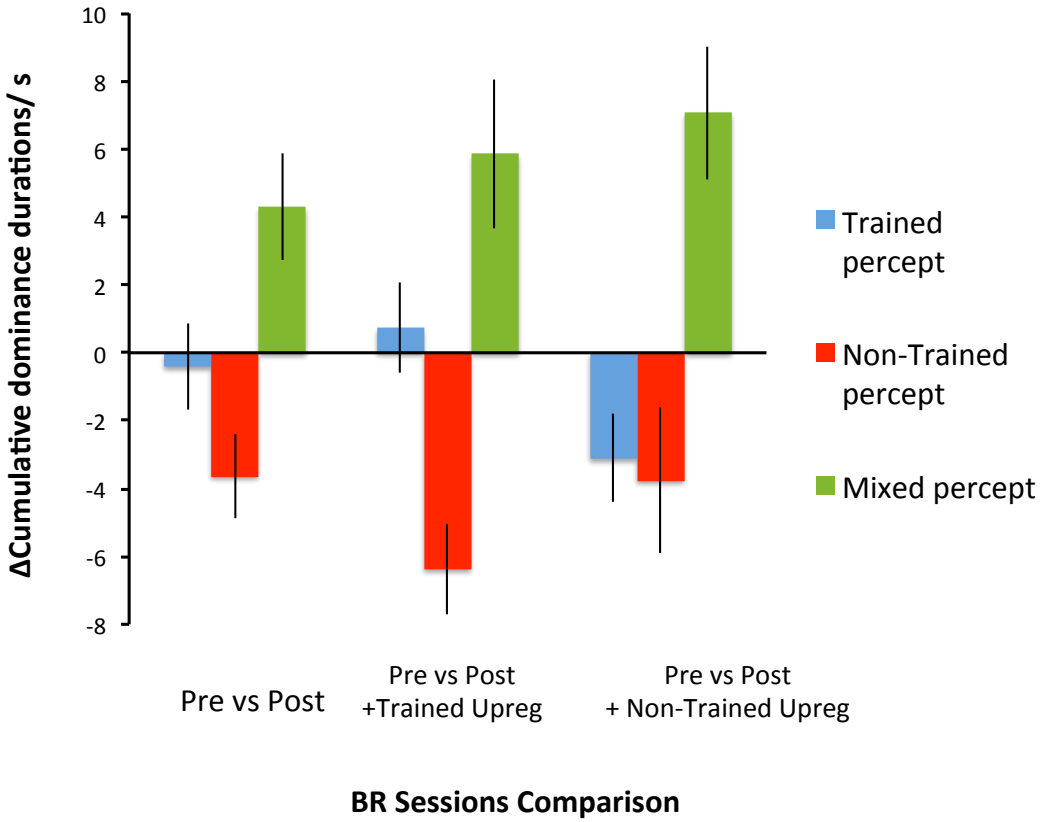

Fig 6-3. Binocular rivalry cumulative dominance durations – differences between conditions.

The difference between the cumulative dominance durations (i.e. the total duration a training strategy related percept, non-training strategy related percept or mixed percept was perceived) is shown above, comparing:-

1. **pre-training baseline BR with post-training BR in the absence of concurrent brain regulation (left);**
2. **pre-training baseline BR with post-training BR with concurrent up-regulation using the trained strategy (middle);**
3. **pre-training baseline BR with post-training BR with concurrent up-regulation using non-trained strategy (right).**

Following training there was a significant reduction in the duration of the percept which was not linked to training i.e. the 'non-trained percept'. There

was also an increase in the duration of the 'mixed percept'. This effect was increased when BR was performed with concurrent up-regulation using the trained strategy. When the non-trained strategy was used, there was an increase in the duration of the mixed percept only (right-hand side). Error bars represent standard error of the mean.

6.4.2 Switches

Planned t-tests were performed for the difference in cumulative switch frequencies, examining pre- vs. post-training changes. Switch frequencies were calculated to specific percepts (grouped in to training strategy related percept, non-training strategy related percept, and mixed percept).

First, the switch frequency for baseline BR was examined pre- versus post-training. There was no significant change in switch frequency between the two time points, for any of the trained percepts (all t's ≤ 1). (See Fig. 6-4 and 6-5).

Comparing post-training BR + trained up-regulation, with pre-training baseline BR, there was a significant reduction in the cumulative switch frequency towards the non-training strategy related percept ($t(9)= 3.86$, $p= 0.003$) . No change was evident for cumulative switch frequency towards the mixed percept ($t(9)=2.08$, $p= 0.07$), nor the training strategy related percept ($t<1$). (See Fig. 6-4 and 6-5).

Comparing post-training BR + non-trained up-regulation with pre-training baseline BR, there was a significant decrease in the cumulative switch frequency towards the training strategy related percept ($t(9)= 4.20$, $p= 0.002$). No change was evident for cumulative switch frequency towards the the non-training strategy related percept ($t(9)= 1.48$, $p= 0.17$), nor the mixed percept ($t(9)= 2.20$, $p=0.06$). (See Figs. 6-4 and 6-5).

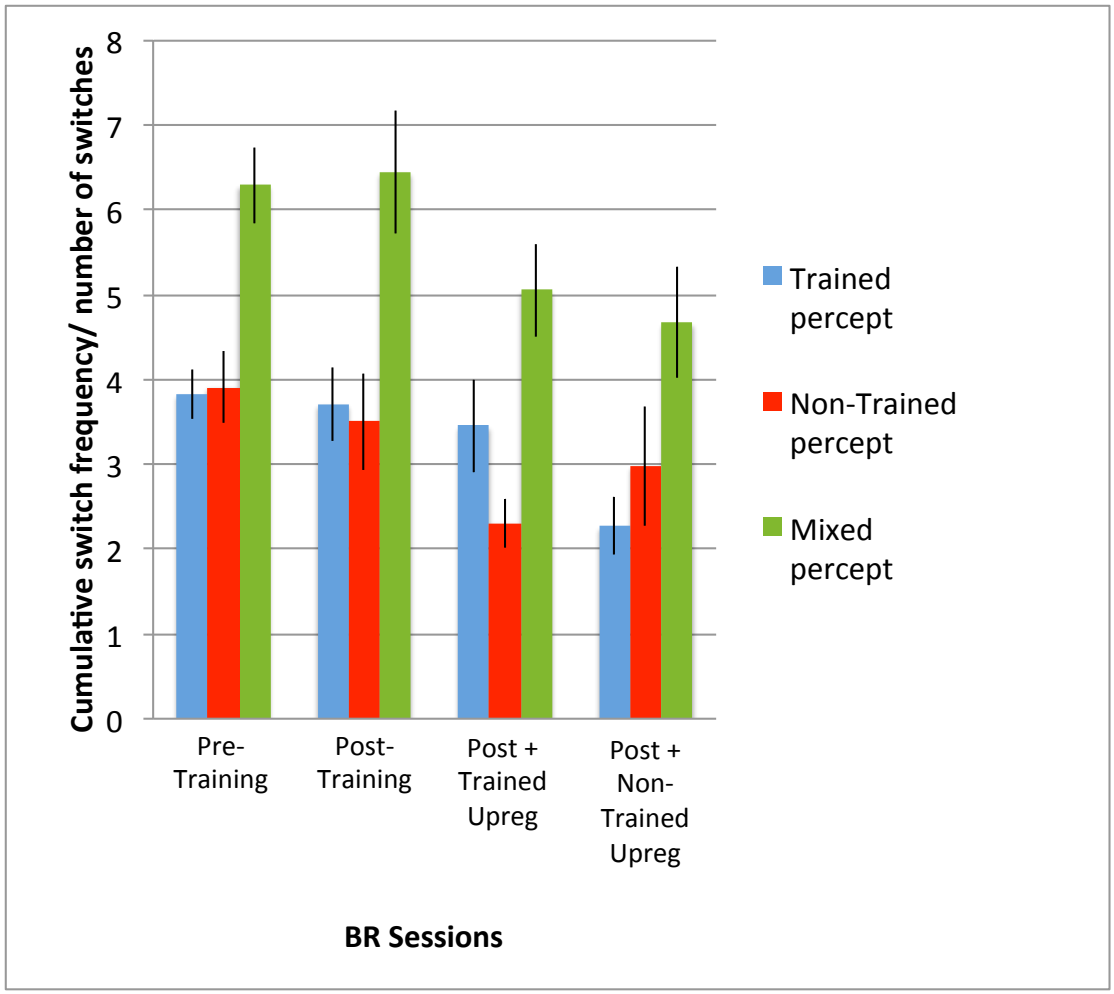

Fig 6-4. Binocular rivalry switch frequencies across participants. Cumulative switch frequencies across participants for pre-training baseline BR, and the three post-training BR sessions i.e. Post training baseline BR, Post training BR with concurrent trained up-regulation, and Post training BR with concurrent non-trained up-regulation. Error bars represent ± 1 SEM.

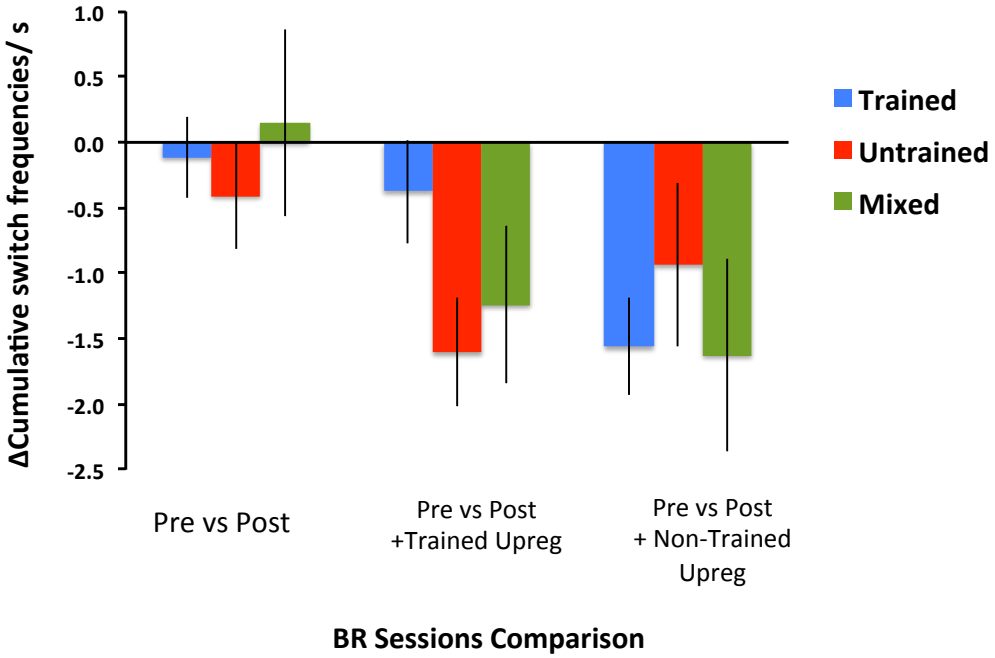

Fig 6-5. Binocular rivalry cumulative switch frequencies – differences between conditions. The difference between the cumulative switch frequencies (i.e. the frequency of switches to a percept) is shown above, comparing pre-training baseline BR with post-training baseline BR in the absence of concurrent brain regulation; pre-training baseline BR with post-training BR + concurrent up-regulation using the trained strategy; pre-training baseline BR with post-training BR with concurrent up-regulation using non-trained strategy. Error bars represent ± 1 SEM.

6.5 Results - Imaging

6.5.1 Neurofeedback training- FFA and PPA

A one-way ANOVA was performed on the 'training' signal (percentage BOLD signal change calculated as the difference between the training-strategy related ROI and the non-training strategy related ROI) across the ten participants over nine training sessions. This indicated a linear increase in the differential percentage signal change over the sessions ($F(1,8) = 5.75$, $p = 0.043$). (See Fig. 6-6).

Planned post-hoc paired t-tests to test whether the training signal significantly differed from zero over the three days were performed. That is, did the BOLD signals from the two ROIs significantly change relative to each other as a function of training, indicating an effect of learning? These analyses showed that there was such a learning effect from Day 2 onwards (Day 1: $t(9)=0.40$, $p= 0.88$; Day 2: $t(9)=3.27$ $p=0.01$; Day 3: $t(8)=2.69$, $p=0.04$). Comparisons were also performed between training day 2 and 3 with training day 1: (Day 2 vs. Day 1: $t(9)= 2.10$, $p =0.06$, Day 3 vs. Day 1: $t(8)= 2.76$, $p= 0.02$, one-tailed t-tests). (See Fig. 6-6).

The learning effect was also examined on a session-by-session basis. There was a significant change in the differential percentage BOLD signal change from training session four onwards [Run 1 $t(9)=1.9$ $p=0.09$; Run 2 $t(9)=0.3$ $p=0.74$; Run 3 $t(9)=0.6$ $p=0.5$; Run 4 $t(9)=3.5$ $p=0.006$; Run 5 $t(9)=3.5$ $p=0.005$; Run 6 $t(9)=2.3$ $p=0.05$; Run 7 $t(8)=2.1$ $p=0.07$; Run 8 $t(8)=2.6$ $p=0.03$; Run 9 $t(8)=3.3$ $p=0.01$].

6.5.2 Transfer session- FFA and PPA

There was strong a priori hypothesis that participants would learn to up-regulate differential activity in the two trained brain regions, in the absence of a neurofeedback signal. This was confirmed by the transfer session following training. During the transfer session, participants performed two additional tasks; BR with

upregulation of the trained strategy and BR with upregulation of the non-trained strategy. Across participants there was an increase in the differential BOLD signal for both of these sessions, which was significant for the first session (t(9)=3.2, p=0.01), but not the second session (t(9)=1.7, p=0.12). (See Fig. 6-6). Overall, the transfer session across participants demonstrated a significant increase in the training signal (t(9)=2.4, p=0.04), and a significant change in brain activity as compared to the first training session (t(9)=2.40, p=0.03, one-tailed).

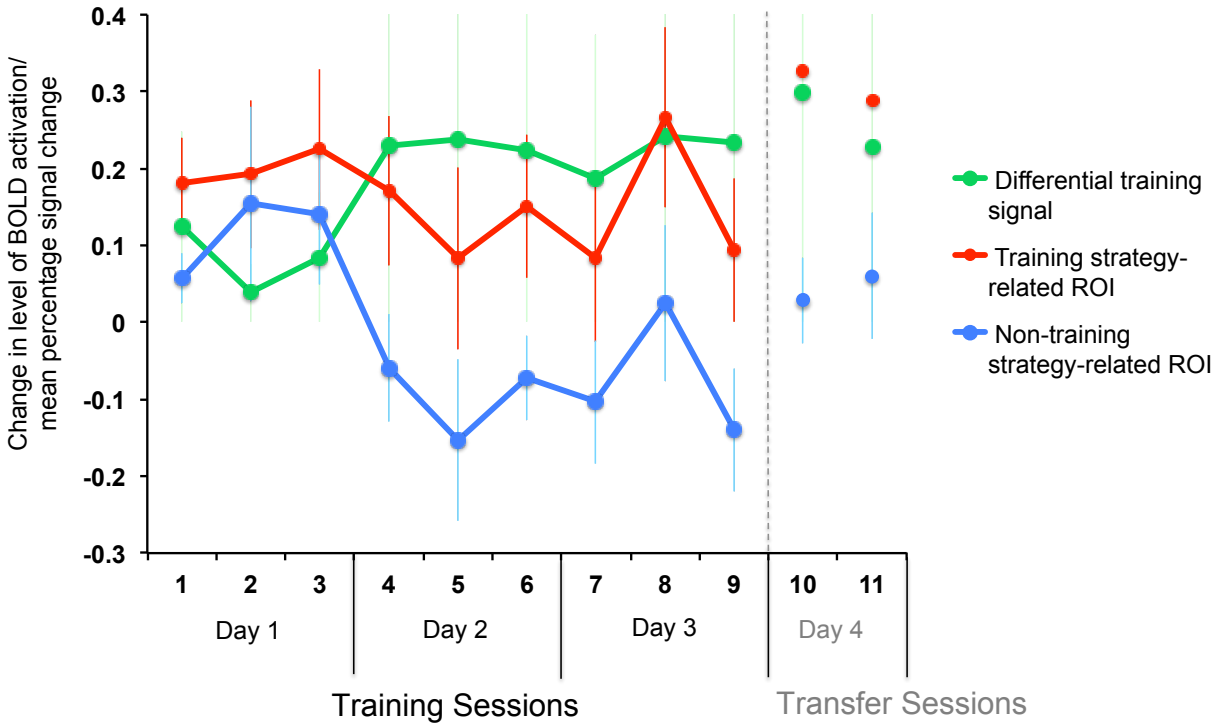

Fig 6-6. Mean BOLD activation changes across groups during neurofeedback training. Activations in the training strategy-related brain region are shown by the red line, and activations in the non-training strategy-related brain region are shown by the blue line, over the nine training sessions, and two transfer sessions. The green line shows the difference in mean BOLD activation between the two brain regions corresponding to the 'training' signal the participants visualised and trained on in the scanner i.e. a 'differential' signal. There was significant increase in this differential brain activation from the 4th training session onwards. Error bars represent ± 1 SEM.

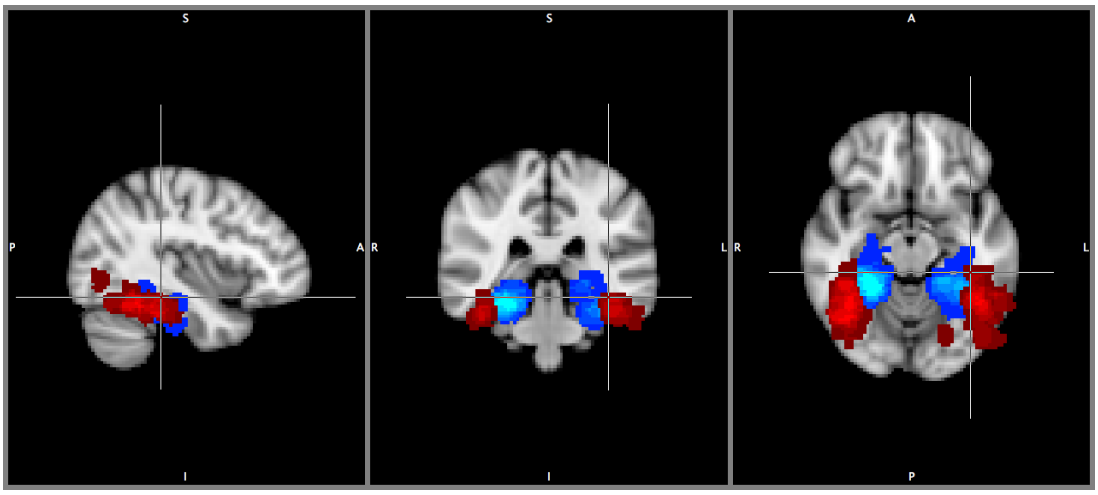

Fig 6-7. Functionally defined ROIs used for neurofeedback training, showing overlap of activation across all participants. The FFA ROI (red) and the PPA ROI (blue) are shown overlaid on a MNI brain.

6.5.3 Binocular rivalry – FFA and PPA

For all of the BR sessions, before and after training, I compared mean BOLD PSC in the training strategy-related ROI, as well as the non-training strategy-related ROI. Across participants, for BR performed before and after training, a significant reduction in activity was seen in both ROIs (Pre BR vs. Post BR comparison: training strategy-related ROI: $t(9)= 3.43$, $p= 0.007$; non- training strategy-related ROI: $t(9)= 2.26$,$p=0.049$); there was no difference in the levels of activity between the two ROIs: $t(9)=0.45$,$p=0.67$). When BR was performed with up-regulation of the training-strategy related ROI, there was a significant increase in the differential brain activation between the training strategy-related ROI, and the non-training strategy-related ROI, compared to zero. A one-tailed t-test was used, as there was a prior hypothesis that the training signal (i.e. based on the difference in activation between the two ROIs) would increase during trained up-regulation, as compared to the non-upregulation conditions ($t(9)= 2.31$, $p= 0.046$, one-tailed). This was underpinned by a significant increase in the level of activation in the training strategy-related ROI (post BR + trained up-regulation vs. post BR baseline: $t(9)= 2.21$, $p= 0.05$). (See Fig. 6-7).

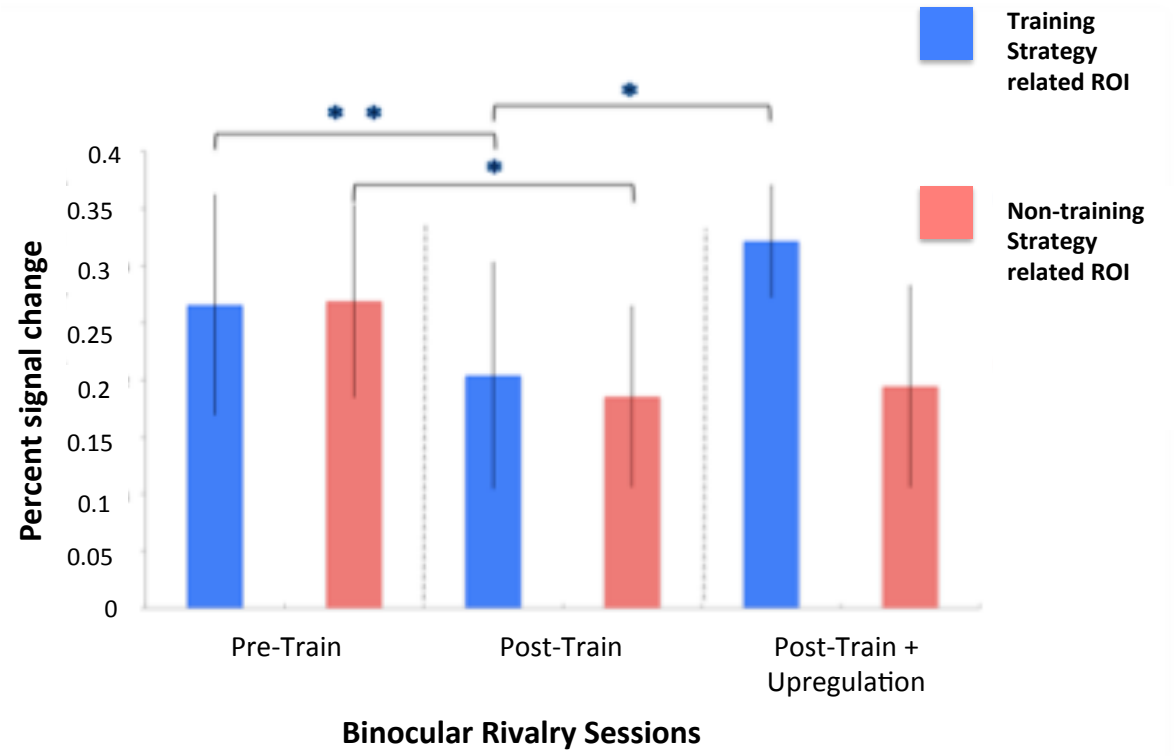

Fig 6-8. Mean BOLD activation changes during BR sessions performed before and after neurofeedback training. The plot summarises the BOLD changes (mean PSC) in trained regions during BR sessions performed before and after neurofeedback training. There was a significant reduction in activation in both the training strategy related regions and the non-training strategy related regions following training (stars indicate significant comparisons). When BR was performed with concurrent trained up-regulation, there was an increase in BOLD activation in the training strategy related ROI only. Error bars represent ± 1 SEM.

When the participants underwent BR, and attempted to up-regulate the non-trained ROI (using strategies which were not specifically trained), no significant changes in brain activity were observed, as compared to the pre- training and post-training baseline BR (all t's< 0.5, one-tailed). (See Fig. 6-7).

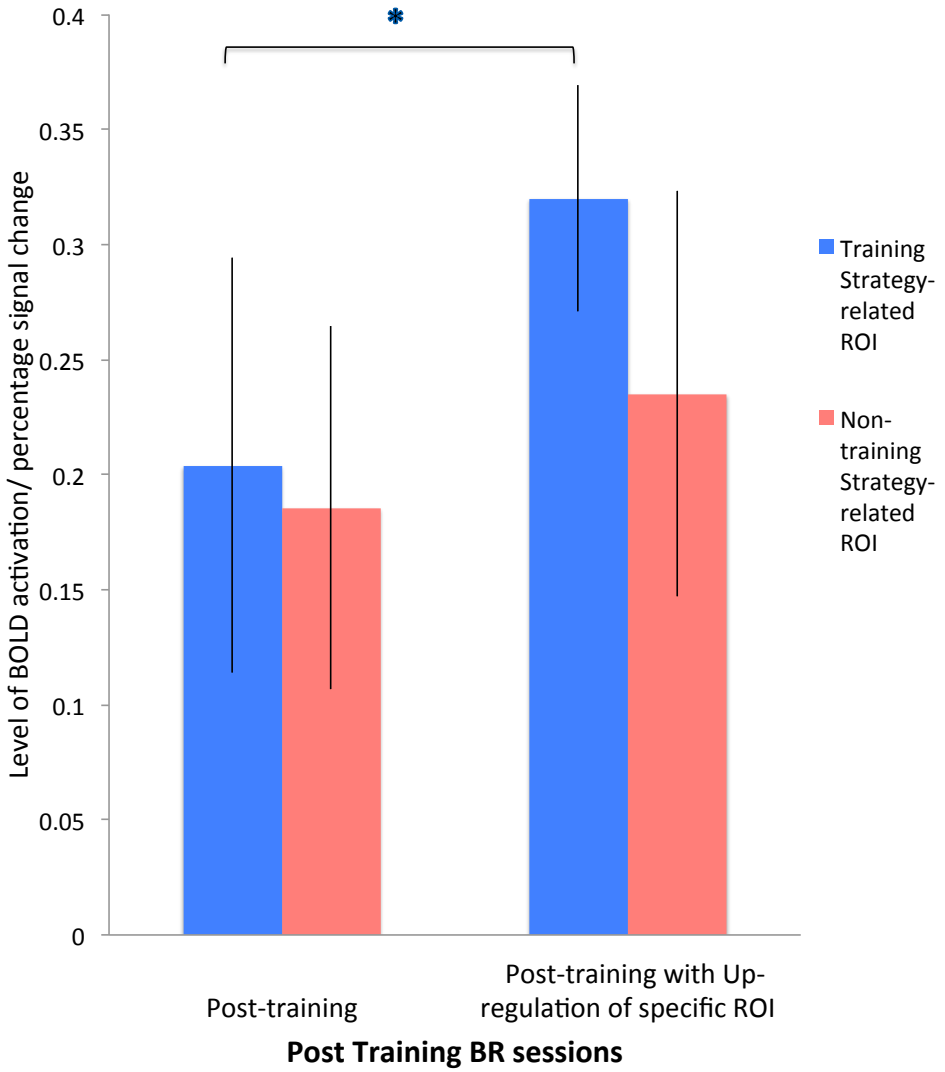

Fig 6-9. The mean percentage signal change in the training strategy-related ROI and the non-training strategy-related ROI following training, comparing the effect of voluntary up-regulation. There was a significant increase in the mean BOLD percentage signal change in the training strategy-related ROI only. (stars indicated significant comparisons). Error bars represent ± 1 SEM.

6.5.4 Neurofeedback training - SPL

As discussed earlier, I specifically examined a SPL for associated changes in activation during neurofeedback training. The co-ordinates selected were based on recent work which confirmed the role of a specific aSPL region during neurofeedback training (81,16).

A second level analysis was performed using a 10 mm sphere small volume correction (x= 22, y= -58, z= 63). This confirmed that there was significant activation in this region during the 9 training sessions (t(9)= 6.45,p= 0.014). BOLD percentage signal changes from this ROI were extracted for the 9 training sessions, and the 2 transfer sessions and plotted as a function of training session (Fig. 6-9).

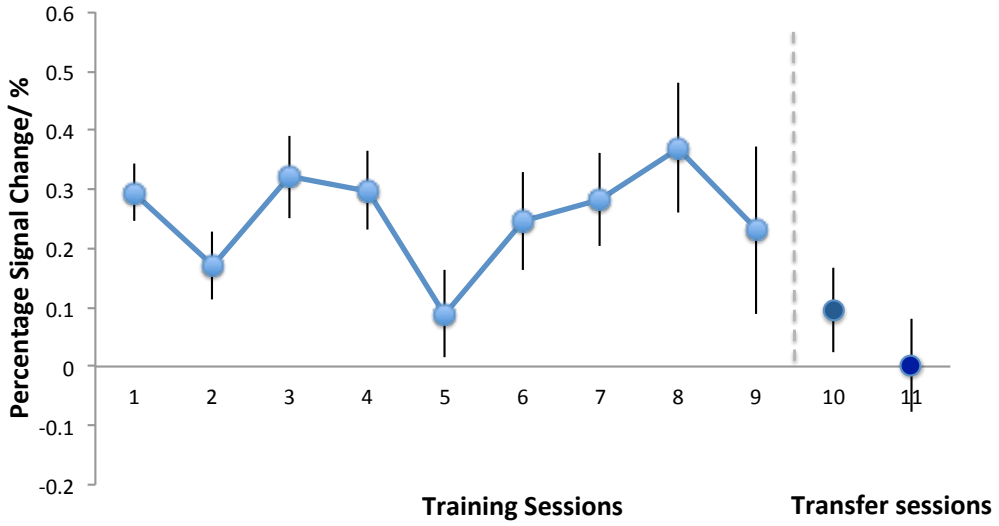

Fig 6-10. Activation curve for SPL region previously indicated as being involved in neurofeedback learning, across 9 training sessions and subsequent 2 transfer sessions, averaged across all participants. Error bars represent ± 1 SEM.

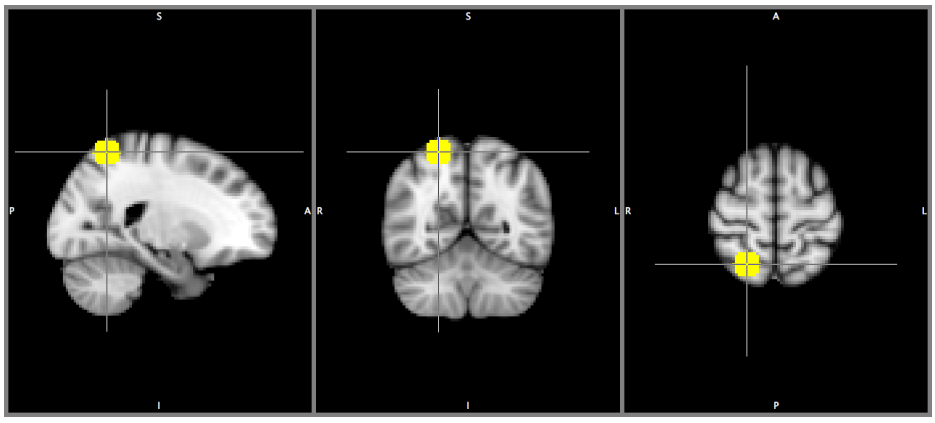

Fig 6-11. Anatomical location of aSPL sphere. View of the 10mm sphere centred on the MNI coordinates 22, -58, 63, located in the anterior SPL.

6.5.5 Transfer session - SPL

For the transfer sessions, the first transfer session produced significant activity in the SPL region ($t(9)= 2.96$, $p=0.01$), whereas the second transfer session did not ($t(9)= 0.60$, $p=0.56$).

Taken together, SPL was active during neurofeedback training while active learning was taking place, with a reduction in activity during the transfer sessions, when learning had been completed and was no longer occurring.

6.5.6 Binocular Rivalry - SPL

Across participants, mean PSC was compared using small volume corrections for the two right sided SPL ROIs 4 mm spheres) previously documented as being involved in BR. These regions were in anterior SPL (aSPL, x= 36, y= -45, z= 51), and posterior SPL (pSPL, x=38, y=-64, z= 32) respectively. Each of the BR sessions before and after training were examined, with and without concurrent trained up-regulation of the training signal.

Pre Training:

During BR performed prior to neurofeedback training, activation was found to be significant in both SPL regions i.e. aSPL, $t(9)=15.89$, $p<0.001$; pSPL: $t(9)=6,43$, $p<0.031$.

Post training

During BR performed following neurofeedback training, activation was found to be significant in both SPL regions i.e. aSPL, $t(9)=12.95$, $p<0.001$; pSPL, $t(9)=6,43$, $p=0.003$.

Post training with trained up-regulation

When BR was performed with concurrent trained up-regulation, following neurofeedback training, activation was found to be significant only in aSPL ($t(9)=19.56$, $p<0.001$).

Pre training BR minus Post training BR

Examining the effects of training on BR, using a comparison of conditions, demonstrated significant activation in aSPL ($t(9)= 7.04$, $p= 0.012$) and pSPL $t(9)= 5.74$, $p=0.03$.

6.6. Results - structural

I examined longitudinal structural changes in FFA and PPA, and SPL using 4mm sphere masks based on the regions of greatest functional activation during neurofeedback training. An ANOVA did not reveal any significant changes in these regions (all p's more than 0.05).

6.7 Results -DCM for neurofeedback training

Prior to training, there was strong bottom-up effective connectivity from the training-strategy related ROI ('St') to SPL and from the non-training strategy related ROI ('Non St') to SPL. These parameter estimates, which correspond to the matrix 'A' in Equation 2, are reported in Fig 6-11 (grey text).

Post-hoc model reduction eliminated all but four parameters in relation to the modulatory effects of training (posterior probability >0.95). The connection from St to SPL was increased by training, denoting increased feedforward connectivity. The connection from Non St to SPL was reduced. Finally, the intrinsic self-connection of SPL was increased (see Fig. 6-10, green text).

Training therefore acted on the feed-forward connectivity of the two trained regions with SPL.

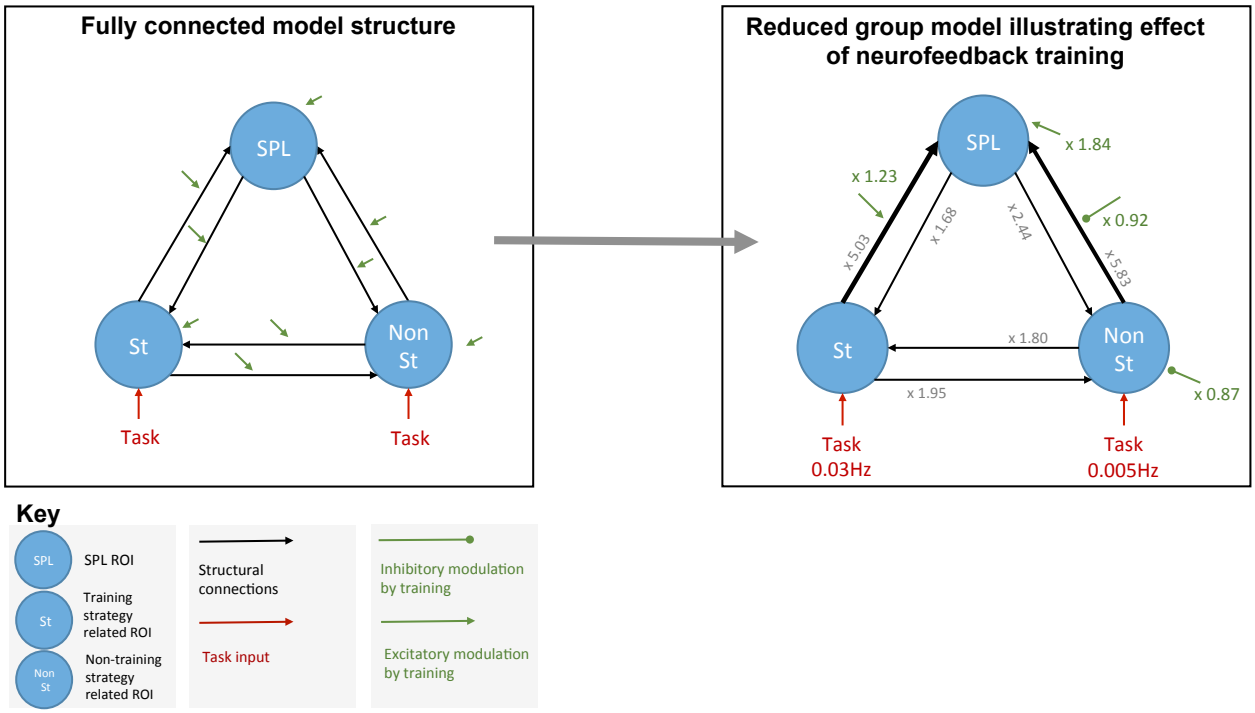

Fig 6-12. DCM results for effects of neurofeedback training. Post-hoc model reduction shows the training effect to be in the bottom-up connections entering SPL, and its self-connections. Excitatory modulation of the Training strategy related ROI to SPL connection, and inhibitory modulation of the Non-training strategy related ROI to SPL connection was observed. Results thresholded at posterior probability 0.95.

The effect of training on the feed-forward connections was further examined, specifically to test whether there was a significant difference between the Training strategy related ROI to SPL connection as compared to the Non-training strategy related ROI to SPL connection. To compare the strength of connection parameters in the group-average DCM, I computed a Bayesian contrast. To illustrate this method, to assess if the first parameter in the model is larger (more positive) than the third parameter, the contrast, c, is a vector:

$$c = [1, 0, -1]$$

The posterior mean of the contrast is calculated by multiplying c by a vector E, which contains the estimated mean of each parameter:

$$\mu = c^{'} E$$

The posterior variance of the contrast is calculated similarly, using estimated covariance matrix V from the DCM:

$$\sigma^2 = c^{'} Vc$$

To report the result, I plotted the probability density function $N(\mu, \sigma^2)$, which shows the probability of a difference in effective connection strengths (a Bayesian contrast), see Figure 6-11. The probability of observing a value greater than zero under this distribution, calculated using the normal cumulative distribution function (CDF), was 97.4%.

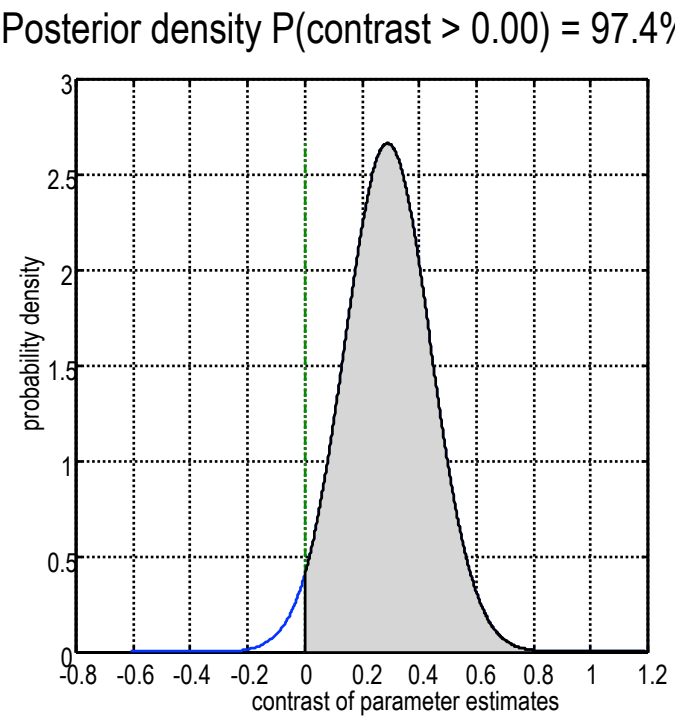

Fig 6-13. Probability that the effective connection strength between the training strategy related ROI and SPL was greater than the effective feedforward connection strength between non-training strategy related *ROI* and SPL.

6.7 Results – Canonical Variate analysis

Plots for specific comparisons of individual differences in behaviour (dominance durations) and functional changes (BOLD changes across training in FFA, PPA, SPL), and behaviour and structural changes (measure of the volume changes in FFA, PPA and SPL following training) are presented below, together with Bartlett's approximate chi-squared statistic for Wilks' Lambda and its p-value, for each comparison.

Two comparisons were examined with regards to the behavioural measures. The first was to examine the change in cumulative dominance durations measured in pre- versus post-training BR. There was no significant association between an individual's dominance duration as a result of training, and the magnitude of their BOLD response in their ROIs ($p = 0.08$, $r = 0.91$; Fig.16.14A). There was, however, a significant association between an individual's behavioural training effect and structural changes in their ROIs ($p = 0.03$, $r = 0.89$; Fig.6.14C).

The next comparison was to assess the relationship between cumulative dominance duration changes (versus pre-training) now during concurrent up-regulation of trained brain activity. The relationship between this behavioural measure and an individual's BOLD response in their ROIs was not significant ($p = 0.05$, $r=0.95$; Fig.16.14B). The same behavioural measure did, however, correlate significantly with structural changes in these regions ($p = 0.003$, $r=0.94$; Fig.16.4D).

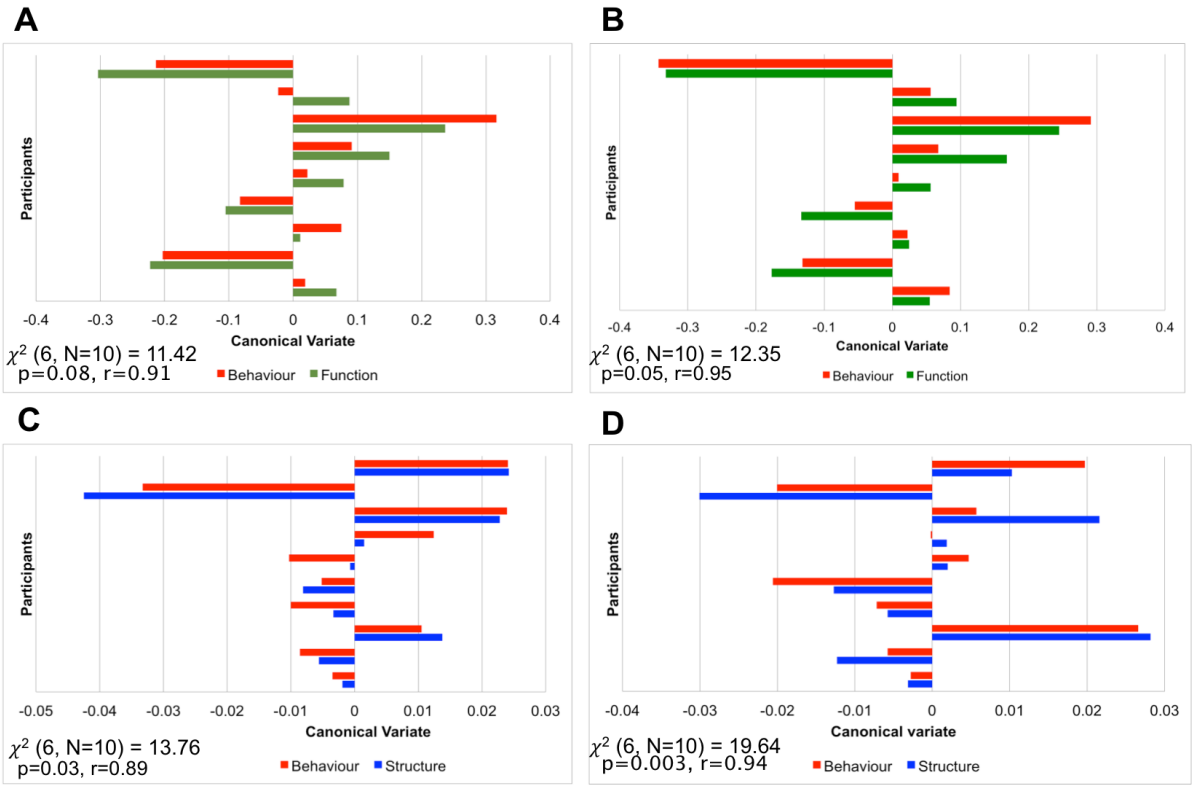

Fig 6-14. Canonical Variate Analysis results. Each of the comparisons/graphs are explained below:

A. Comparison of behaviour (i.e. change in cumulative dominance durations for training strategy related percept, non-training strategy related percept and mixed percepts between pre-training BR and post-training BR) and function (i.e. change in the level of the differential BOLD signal between training strategy related brain region and non-training strategy related brain region).

B. Comparison of behaviour (i.e. change in cumulative dominance durations for training strategy related percept, non-training strategy related percept and mixed percepts between pre-training BR and post-training BR performed with concurrent trained up-regulation of brain activity) and function (i.e. change in the level of the differential BOLD signal between training strategy related brain region and non-training strategy related brain region).

C. Comparison of behaviour (i.e. change in cumulative dominance durations for training strategy related percept, non-training strategy related percept and mixed percepts between pre-training BR and post-training BR) and structural changes after training (in trained brain regions FFA, PPA and SPL).

D. Comparison of behaviour (i.e. change in cumulative dominance durations for training strategy related percept, non-training strategy related percept and mixed percepts between pre-training BR and post-training BR performed with concurrent trained up-regulation of brain activity) and structural changes after training (in trained brain regions FFA, PPA and SPL).

6.8 Discussion

This study demonstrates that neurofeedback training produces lasting changes in bistable visual perception as measured with BR. Using fMRI with neurofeedback, participants were trained to differentially modulate the level of activation in two higher-order, category-specific visual brain regions, FFA and PPA. These regions have recognised roles in the processing of perception related to face stimuli and house stimuli, respectively. Following training, sustained changes in the perception of stimuli associated with these regions, were found during BR. These changes in perception were underpinned by functional changes in actively trained brain regions, as well a higher order control-related brain region, aSPL during neurofeedback training. DCM analyses showed changes in effective connectivity between aSPL and the trained brain regions during neurofeedback training, implicating the interplay between a control-related area, and higher-order visual areas involved in BR. A CVA analysis confirmed a significant association between changes in BR dynamics (i.e. dominance duration) following neurofeedback training, and structural changes in FFA, PPA and SPL, following neurofeedback training.

To summarise my experimental findings, participants learned to up-regulate the neurofeedback signal in both groups (i.e. 'Face group', 'House group'). Both groups were unaware of the underlying nature of the training signal. From the five potential means by which the training signal could be increased (see page 221), participants *decreased* the level of activity in the non-training strategy-related ROI, while appearing to maintain or 'clamp' the level of activity in the training strategy-related ROI. Comparing this to observed behavioural changes in BR dynamics, before and after training, there was a significant reduction in the cumulative dominance durations of the non-training strategy related percept. This was in keeping with the direction of the brain activation changes over training. The behavioural change was further enhanced when BR was performed in the presence of concurrent up-regulation of the training- strategy related ROI.

Examining the imaging findings during BR, specifically the concurrent BOLD changes during up-regulation in training and non-training strategy-related ROIs,

there was a significant increase in the level of activity in the training strategy-related ROI only, which is an intuitive result. However, training produced a different effect on neural activation in the ROIs from that observed during BR. During neurofeedback training, the difference between the two regions was driven by a significant *reduction* in the level of activity of the non-training strategy-related ROI, while during BR, the participants were able to drive the level of activation in the training strategy-related ROI. An explanation for this maybe related to the presentation of visual stimuli during BR. As one of the presented stimuli was salient to the training strategy (e.g., face stimulus for the 'face' group) and involved a trained category-specific ROI (e.g., FFA), concurrent up-regulation in the presence of trained, salient stimuli may have resulted in the enhanced response in trained category-specific cortex. In contrast to this, no directly salient visual stimuli were presented during the neurofeedback sessions, with neural activations in the trained brain regions being primarily self-generated by the participants.

The results from my study may imply a causal relationship between the modulation of brain activation and changes in perception, and builds on previous work illustrating that neurofeedback training of cortical regions using a relatively unconstrained measure of activity such as mean regional activity, produces changes in the function of the region. I found that modulation of the relative level of activation between two higher-order visual brain regions using neurofeedback training altered the perception of stimuli linked to the trained regions.

In a related fashion, Shibata et al. (509) demonstrated lasting changes in the ability of participants to detect particular patterns (orientations of gratings) following neurofeedback training of primary visual areas. Scharnowski et al. (16) on the other hand, only demonstrated improvement in a visual detection task involving grating orientation detection, when concurrent trained up-regulation was performed, confirming an effect of learning to regulate brain activity. These two studies illustrate an important point – neurofeedback training affects the activity and/or functionality of the brain region.

Neurofeedback training has been conceptualised as a specific form of instrumental conditioning that utilises a BOLD signal based training signal as the operant value

(86) Successful control of the signal produces an implicit reward, and is selected such that brain activation in target brain region/s is directly linked to the task at hand. In order to modulate the activity in these brain regions in the absence of an external stimulus, participants must utilise abstract cognitive strategies to voluntarily regulate the on-going activity within a specific brain region; frequently these strategies use mental imagery (85). Studies have previously demonstrated participants gaining control of the feedback signal in the absence of an explicit instruction (232,510,511), although this has been countered (511). Scharnowski and colleagues (16) altered their experimental paradigm during their study to include explicit instruction related to the strategy to be used during neurofeedback. They demonstrated a more rapid learning effect as a consequence. In my study the training signal was produced by differential activity between two brain regions. However, explicit instruction for control of the signal was related to only one of the two brain regions. The second brain region acted as an internal 'control' (329) in relation to non-specific global changes in brain metabolism and hence neuronal activity e.g. related to arousal, attention, breathing (57,58,301,315). Interestingly, although the participants were not offered a strategy that required involvement of the secondary brain region, they were nonetheless able to control its activity. This provides an interesting example of separating brain activity in a specific cortical region from its functional output.

More specifically, I examined participants' ability to regulate the non-training strategy-related region as a separate condition by providing an explicit instruction linked to that region, but without any training. Participants were unable to modulate the training signal with a non-trained strategy (see Fig. 6-7), despite having previously successfully modulated the activation levels in both ROIs, when using a training strategy linked to only one of the two ROIs. This observation identifies the importance of the neurofeedback training process, in that control of the visual interface/ 'training' signal is qualitatively different to a one-to-one mapping of mental imagery and activation of imagery-content related brain regions. Further, it illustrates the specificity of the neurofeedback training process, implicating a procedural component to the underlying process (86,94). In terms of the behavioural changes in BR, during attempted use of non-trained mental imagery strategies, participants demonstrated a significant increase in the duration of the mixed percept, but not the non-training strategy related percept.

The mechanisms underlying the neurofeedback effect are unknown, and the literature examining this topic limited (72,77,332,334,512–514). The complexity of the neurofeedback process relates to a number of observations. To begin with, neurofeedback-driven brain modulation potentially provides access to neural processing which are normally hidden. The issue of causality emerges - mentalising modulates pre-existing neural processing, shaped by the specificity of controlling the visual interface. Several neural processes are likely to be involved, including the binding of neural and mental events (515), agency (516), introspection (517), attention and neuroprosthetic control (94). Top-down co-ordination is therefore likely to be necessary during neurofeedback control, with a network of regions being engaged during learning. Midline structures including the anterior and posterior cingulate have been suggested (77). More specifically in the context of visual perception, Scharnowski et al. (16) observed increased top-down connectivity between aSPL and primary visual cortex in participants who successfully learned to up-regulate activation in the visual cortex. SPL has a recognised role in the top-down control of covert visuospatial attention and in cognitive control (454,518). Using an aSPL ROI that has been previously implicated in neurofeedback (81), I found significant activation during neurofeedback training in the present study.

A DCM analysis was used to investigate changes in effective connectivity between the trained brain regions which demonstrated significant task-related BOLD activation, and infer the directionality of changes to potentially identify causal relationships between brain regions examined. DCM extends the assessment of statistical dependencies between two brain regions that examine correlations between fMRI time series (i.e. functional connectivity) to include the specification and testing of forward models, which include 'hidden' biophysical and neuronal states i.e. region specific haemodynamic models, of how the observed data occurred (519). By doing this, 'connections' are inferred by parameterising the coupling between brain regions. More specifically, both the haemodynamic and neuronal contributions to the BOLD response are modelled using non-linear differential equations in continuous time (i.e. dynamic), which incorporate the strength of the connections and how they are effected by experimental factors. DCM therefore models the coupling interaction or 'effective connection' between brain regions, which in the case of a putative higher

order brain region interacting with a lower tier brain region i.e. SPL and FFA, can be feedforward or feedback(520). Finally, DCM (See 6.3.5.2 and (504,521)) provides for the possibility of the connection being either excitatory or inhibitory. The effect of neurofeedback training in this study was found to positively modulate the excitatory feedforward connection between the training-strategy related ROI to aSPL, and negatively modulate the excitatory feedforward connection between the non-training strategy related ROI to SOL. There was also a positive modulation on the excitatory self-modulatory connection of aSPL. More complex DCM models could be envisioned, potentially involving other brain regions implicated in cognitive control and perception, such as the dorsolateral prefrontal cortex, as well as other parts of the parietal lobe i.e. posterior parietal cortex (71,454,497,499,518,522). Although this would pose a significant challenge in terms of computational analysis, it may nonetheless enable a more complete view of the way in which perception is controlled, and provide further targets for neurofeedback-led modulation.

The feedforward interaction between the trained regions and SPL may be in keeping with a previously posited theory of neurofeedback learning (90). According to this 'two-step model' of neurofeedback learning, feedforward processes interact with feedback processes, with higher-order processes such as attention being utilised in a trial-and-error search linked to the provided instruction. Once the skill is learned, it becomes an automatic process, and is stored in implicit memory; at this point top-down involvement is no longer required. In my study, in line with this mechanism, activation of SPL was observed during the neurofeedback training process, with increased feedforward connectivity from the directly trained brain regions. Following this period of training, activation in SPL reduced to zero during the transfer sessions.

Turning now to the perceptual task, BR provides a useful experimental paradigm to examine the effects of voluntary brain modulation on perception. It involves the simultaneous presentation of two stimuli, separately to each eye. The resulting perceptual phenomena, where perception of one stimuli competes with other, is stochastic, but is amenable to manipulation (479,523–525). A number of mechanisms have been used to explain the neural workings of BR, invoking both top-down, higher-order cognitive modulation, and bottom-up, salience driven processes. Attentional processes have been implicated, and a shared locus for

attention and perceptual bistability has been shown (526). BR has been used to assess targeted modulation of higher-order brain regions, which have explicit roles in visual perception. The addition of neurofeedback training enables a direct means of manipulating this, combining aspects of mental imagery, perceptual biasing and learning.

An alternative account of BR conceptualises it as a predictive or inferential process. It adds to the previously described feed-forward model (i.e. tiered information processing through an increasingly complex cortical hierarchy) by allowing for a reciprocal interaction between the internal mental state and the external stimulus acting upon it (168,497,527). Taking the proposed Bayesian framework, the dominant percept at any given time is maintained by the highest posterior probability, at the top of the cortical hierarchy. The stimulus representation at the lower levels provide error signals which are compared continuously with top-down predictions- the percept is rendered more or less stable in relation to this bottom-up inhibitory mechanism i.e. the lower the error signal, the more stable the percept, the higher the error signal, the greater the suppression of the percept (138,527,528). In Bayesian terms, the initial prediction for this study would be that the a strong prior manipulation by neurofeedback entrainment of category-specific brain regions would cause perceptual stability to shift in the direction of the trained region. More specifically, I observed BR dynamics being shifted in the direction of the information represented by the trained visual brain region. The shift in BR dynamics observed following neurofeedback training was exaggerated during concurrent trained up-regulation. The application of this Bayes-related theory to the observed findings of this study therefore provides one possible unifying explanation.

The behavioural changes during BR after neurofeedback training may arise from a number of related but distinct processes. Mental imagery does play a role in neurofeedback, although as discussed above, it has been shown to be insufficient as an explanation for the behavioural changes. Rademaker and Pearson (486) examined whether successive training of mental imagery resulted in increasing ‘vividness’, and if this had an impact on perceptual rivalry. No effect of training was found in relation to the vividness of the imagery, nor was there any benefit in expending increased effort in terms of imagery. It is therefore unlikely that instructing

the participants to use mental imagery was itself sufficient to produce the behavioural changes observed in this study. Similarly the role of 'priming' and 'cueing' could also be invoked as potential arguments for the observed behavioural changes. Denison et al. (529) showed that prior presentation of a specific orientation grating caused an increase in the perception of the identical grating during BR. However, this did not produce a lasting effect on BR dynamics – dominance durations were not significantly changed in a lasting manner. Exogenous cueing prior to BR may increase the probability of the predominant percept being linked to the cue. Indeed, prior to BR, hearing sentences with the word 'face' result in FFA activation, while hearing sentences with the word 'house' result in PPA activation (515). Nonetheless, Pelekanos et al. (530) did not find any significant change in stimulus dominance between faces and houses on rivalry trials when participants had been cued with a word linked with one of the rivalrous stimuli. Evidently neither altering the level of activity in higher order brain regions involved in perception, nor applying known influences on visual perception, provide a comprehensive explanation for the lasting shifts in perceptual bistability observed following neurofeedback training in this study. My findings, along with others, suggests instead that the interaction of mental imagery and control of a physiological measure of brain function using neurofeedback, can be used to produce lasting changes in behaviour – in my study, a significant change in the unconscious processing of visual stimuli.

The work presented above was performed on a limited number of participants. Nonetheless, a more detailed examination of the individual differences in neurofeedback training, and its effect on perceptual dynamics during BR, functional and structural changes in the trained brain regions, was performed using CVA (507). This is a multivariate analysis technique that can accommodate multiple individual measures such as behaviour, indices of connectivity, structure and function in order to establish the extent to which behavioural changes linked to neurofeedback training of target brain regions, predicts functional and structural changes. The extent to which BR dynamics, specifically the changes in cumulative dominance durations of each of the percepts, was altered by neurofeedback training of FFA and PPA, was found to be linked with measures of structural changes in FFA, PPA and SPL following neurofeedback training. This may suggest that neurofeedback training, even over a relatively short period of time (1 week) can alter perception, as a result

of changes in functional and structural plasticity of trained and linked brain regions (531,532). Further, longitudinal assessments in relation to both lasting training effects, and adaptive changes in brain plasticity could be envisaged with a larger participant cohort.

In terms of the implications of this study, it illustrates the impact of brain-based training approaches, over and above simply task-based training. It provides for the possibility of directly manipulating brain function to alter perception by targeting relevant brain regions rather than behavioural routines. Further, control brain regions could serve as neurofeedback targets to explore the possibility of more efficient neurobiological training paradigms. From a behavioural perspective, it is now to be established the extent of the unconscious biases that can be trained and thereby altered using neurofeedback, such as social and potentially aberrant biases. Finally, the work presented demonstrates an operational platform for non-invasively modulating network brain functioning to produce changes in behaviour and cognition – this can be translated into the clinical arena, for the purposes of modulating functional networks disturbed by disease such as PD.

6.9 Conclusion

To my knowledge, the work presented above provides the first evidence of the effects of neurofeedback training on higher order visual perception. This was implemented with rt-fMRI, to produce unconscious shifts in the perception of higher-order visual stimuli. Trained up-regulation of ventral temporal visual areas produced linked shifts in in the stimulus perception as measured with BR behavioural measures. They were correlated with functional changes in the trained brain regions, as well as changes in effective connectivity with top-down cortical ‘control’ regions implicated in neurofeedback learning. Individual differences implicated structural plasticity in the brain regions examined. Future work can investigate the extent to which neurofeedback training can be used to further alter high order processes such as consciousness, potentially by targeting perceptual and well as control brain regions.

7. GENERAL DISCUSSION AND CONCLUSIONS

7.1 General summary

The experimental studies in this thesis investigated the development of a non-invasive BCI, implemented with rt-fMRI, targeting brain activations in higher order visual areas for the purposes of an attention driven communication interface, and altering perception using neurofeedback. Both of these applications may be considered examples of 'cognitive' BCIs.

In this general discussion, I will review the findings of the experimental chapters, and relate them to the aims of this thesis – namely the extent to which rt-fMRI was successfully used for the practical applications of decoding and neurofeedback. Further, I will discuss the extent to which the results of the studies provide mechanistic and potentially causal insights into visual cognitive functioning. I will consider current and future experimental work that could/are being used to test hypotheses generated from the studies presented here, specifically examining the use of non-invasive BCI in the control of cognition, and behaviour, in health and disease.

The experimental studies can be considered in the context of the decoding of attention (Chapters 3 and 4), applications of rt-fMRI for the purposes of a communication BCI (Chapters 5), and the use of neurofeedback-based training to modulate perception (Chapter 6).

7.2 Decoding of attention

The first two chapters examined the classification of specific patterns of brain activations in higher order visual regions during the deployment of spatial attention.

Initially, using multivariate classification approaches, spatially distributed brain activations evoked in association with covert shifts of attention to simple visual stimuli in one of four quadrants of the visual field, could be classified at above-chance accuracy. The cortical ROI investigated included primary visual cortex, category-specific brain regions with proven retinotopy such as lateral occipital cortex (LOC) (49,112,384), and putative control regions implicated in top-down modulation of attention including the parietal lobe (123,444,449), and FEF (116,121,123). Regions with a hypothetical role in spatial attention such as the SMA were also explored (533,534). Overall, results from these studies provided classification accuracy only marginally (though significantly) above chance, with the most successful classification being obtained from primary visual cortex and LOC. This was specifically true for four-quadrant classification, which was the most important in terms of applicability as an attention-driven BCI (i.e. highest degrees of freedom for control of an attention driven BCI).

This finding was likely to be related to a number of known attributes of the visual system. Primary visual cortex is highly sensitive to the physical properties of a stimulus (e.g. luminance, contrast, location). Further, the magnitude of BOLD activation elicited by visual stimuli, is comparatively less in these regions as compared to higher order visual cortical areas (535). The relative success of multivariate classification therefore confirmed the presence of attention related information in primary retinotopic cortex, albeit at level insufficient to be identified by magnitude-driven univariate approaches. Further, the presence of statistically reliable classification in the LOC may be attributed to the robust retinotopy of this region (112), together with its role in object processing (48,49,452) – the use of spirals may have activated cortical regions sensitive to retinotopic location, as well as object-sensitive areas in higher order cortex.

In this experiment, multivariate classifiers performed significantly better as compared with simpler univariate approaches. The relative simplicity of the stimuli used in the paradigm (spirals), as compared to high contrast gratings or checkerboards (e.g. 14) would potentially mean that multivariate classification would be more likely to identify distributed fine-grained information in the brain activations observed. Conversely, univariate approaches that are primary dependent on the magnitude of derived brain

activations may lack sufficient sensitivity under some cognitive conditions and/or classes of visual stimuli, to reliably reveal stimulus location in primary visual cortex, and stimulus properties in LOC.

Nonetheless, this work provided a number of insights into establishing a robust decoding platform. Building on these observations, in the second study (Chapter 4) I used more complex, real-world visual stimuli. This was done with the intention of taking advantage of category specific information coded in higher-order visual areas in addition to spatial information, in order to enrich the brain signal being classified. I also implemented a novel use of m-sequences for the purposes of quadrant-specific stimulus presentation, thereby adding a further source of information for the decoder. Once again primary visual cortex was examined in relation to classification accuracies. However I now extended the approach to include other category-specific brain regions such as FFA, in addition to LOC, as well as performing a detailed examination of the different parts of the parietal lobe.

Using univariate approaches I found above chance classification for four-quadrant attentional deployment, in higher order visual cortex, including parietal regions, FFA and LOC. Further, the magnitude of classification accuracy was improved as compared to the results obtained in the first study. Multivariate classifiers were investigated again, including more sophisticated approaches. However, univariate classification accuracies were much improved in the second study i.e. i.e. 50% classification accuracy for four-quadrant classification, while multivariate accuracies remained just above chance i.e. 30% classification accuracy for four-quadrant classification. The improvement in classification using univariate approaches in the second experimental study may be attributable to the strength of BOLD activation produced by using more complex visual objects, focused within a specific cortical location. Further, I used a winner-take all decision algorithm for the univariate classification that identified the highest level of activations produced in relation to the deployment of spatial attention to one of four potential streams of stimuli in each quadrants. As this univariate approach produced markedly improved classification accuracies for the deployment of spatial attention, I went on to develop it for online implementation with rt-fMRI.

7.3 Online decoding with rt-fMRI

After establishing a reliable decoding pipeline for the classification of spatial attention using brain activations linked to category-specific information, quadrant- specific stimulus presentations using m-sequences, and spatial location, I sought to replicate this using an online rt-fMRI platform. An analogous paradigm was successfully implemented but now using rt-fMRI, with four-quadrant classification performed on a trial-by-trial basis. Overall classification accuracies computed in real time were comparable to those calculated offline from LOC, FFA and parietal lobe, with individual decoding accuracies reaching rates of 70%. It was also shown that decoding accuracies were highest during the first half of each experimental session, and these correlated with more rapid reaction times on a n-back task. As such, I concluded that within-session fatigue may have been a factor affecting classification accuracy, and that this would potentially need to be accounted for when implementing a BCI. Specifically, this would be of particular importance in a clinical context, for example with physiologically frail patient groups e.g. locked-in-syndrome, when attempting use of an attention driven BCI communications platforms (15,86,536).

7.4 Neurofeedback

In the final experimental chapter, a rt-fMRI neurofeedback loop was established to examine the modulation of activation in ventral category-specific brain regions involved in visual processing. In order to test if neurofeedback training altered the function of these brain regions and as a consequence, changed visual perception, I used a BR paradigm, with the presentation of visual stimuli, whose processing was linked to the brain regions being trained. The regions trained were PPA and FFA, with participants learning to modulate the difference in activation levels between the two brain regions. The ‘face’ group, that is, those participants who learned to increase the levels of activation in FFA relative to PPA, perceived house stimuli less than face stimuli i.e. reduced cumulative dominance durations of house stimuli. Similarly, the ‘house’ group, that is, those participants who learned to increase the

levels of activation in PPA relative to FFA, perceived face stimuli less than house stimuli i.e. reduced cumulative dominance durations of face stimuli.

This might suggest that by increasing the level of activation within functional brain regions involved in the processing of category-specific visual stimuli using neurofeedback training, neuronal representations relying on the function of these regions were strengthened. This may in turn have lead to more rapid processing of these stimuli, and as a result, perceptual biasing. A conclusion from this experiment, was that these findings constituted a neural analogue of Levelt's 2nd proposition (490). The original proposition, which has since been refined (537) was based on the physical properties of visual stimuli and its psychophysical relations with BR alternation dynamics. It states the following:

"Variation of the stimulus strength in one eye will only influence the mean dominance duration of the contralateral eye and not the mean dominance duration of the ipsilateral eye."

Levelt's historic observations have informed computational (527,538) and cognitive accounts of BR (539). In the context of the experimental findings in Chapter 6, it may be useful to consider a neural formulation of Levelt's 2nd proposition, with regards to perceptual biasing. Specifically, strengthening the neuronal representation related to specific visual stimuli may bias perceptual processing, such that stimuli linked to these representations are more likely to be preferentially processed, resulting in unconscious perceptual biasing. A further qualification in relation to intrinsic vs. extrinsically manipulated brain activations might be that in the absence of direct stimulation through repetitive exposure, modulation of neuronal representations may only be achieved by intrinsically stimulating brain activation.

The investigation of the learning process underpinning the observed effects of neurofeedback training led to the investigation of a possible control region, SPL, which had previously been implicated in other types of learning involving the visual system (540). This was regarded as being important from a number of perspectives; the identification of such a region may provide a more direct target for the facilitation of learning in neurofeedback training paradigms. Future training paradigms could

seek to specifically train connectivity parameters between such regions and the specific regions linked to the process of interest (e.g. *24*) . By engaging these regions more directly, a more rapid and robust learning procedure may be facilitated.

With regards to understanding the mechanism underlying neurofeedback-based learning, a DCM connectivity analysis focusing on the trained regions as well as putative control regions provided an opportunity to investigate the direction of information flow (e.g. *25*). The results of the DCM analysis indicated a two-step process, with training initially increasing feedforward connectivity between the trained regions and SPL, followed by a reduction in activation in SPL after the active learning phase. This result is in keeping with a previous conceptual proposal made by Lacroix (90) outlining a dual-step process for biofeedback learning, although there remains a need to integrate this with a procedural understanding of how neurofeedback training leads to learned regulation of brain activation (e.g. *27*), and how such training gains might be sustained.

I examined structural changes in relation to neurofeedback training. No significant results were obtained at the group level. However, a multivariate combinatorial analysis of all of the measures of neurofeedback training using CVA, demonstrated individual differences in the relationship between the behavioural changes in BR dynamics, and the structural measures. This potentially important finding is suggestive of functional and structural plasticity following 1 week of neurofeedback training with rt-fMRI. Rapid adaptive plasticity as recently discussed in the training literature (233) has not been previously evidenced in relation to structural changes following neurofeedback training (542), although there has been some encouraging data in relation to functional activation changes (13,66).

7.5 BCI technology - Perspectives and insights

I will now consider how the work performed in this thesis fits into the broader context of our understanding of BCI technology. The development of BCI platforms provides an opportunity for cognitive, motor and behavioural enhancement, in health and

disease. Focusing predominantly on the prospects for cognitive enhancement in the context of the work in this thesis, I discuss current and emerging applications under the two headings of ‘software’ and ‘hardware’ upgrades. This is a conceptual simplification borrowed from the computing fields to assist with understanding a complex, and highly interconnected field – as such there remains a significant overlap in the concepts and applications described below.

7.5.1 A ‘software upgrade’: Shaping neural representations using neurofeedback

An emerging implication of my work and that of others is that neurofeedback training, using mental imagery related to the specific functional processing of brain region, may enable ‘shaping’ of neural representations (88), without participants being consciously aware. This neurofeedback driven instrumental conditioning could be applied in sensorimotor cortex, or in perceptual regions. I discuss the importance of mental imagery in relation to accessing neural machinery in a desired cortical region, prior to exploring potential applications of this concept.

Neurofeedback training, particularly for the purposes of cognitive BCIs, utilises mental imagery as a surrogate for engaging neural processes. Mental imagery acts as a form of weak perceptual phenomena, which is modality-specific (543–545). I will focus on empirical evidence relating to this view of visual mental imagery, and of its relevance to this study. Visual imagery produces activation throughout the visual hierarchy, with some important quantitative and qualitative differences depending on the cortical region (535,546). Early visual cortex, at the bottom of the hierarchy, demonstrates an overlap in activation patterns between visual imagery and explicit visual perception. There is a high degree of correspondence in relation to the physical properties being imagined/perceived e.g. retinotopic location, spatial frequency, edges. In comparison, imagery-related activation in higher order cortex, such as the ventral visual areas tend to be invariant to visual detail (35). Activations in these areas are related to the semantic content, and are more flexible and abstract (547). Further, activity patterns are more reliably decoded from higher order

visual areas, leading to the suggestion that perceived and imagined visual images resemble each other more closely in higher order visual areas (535).

Therefore, the use of higher order visual areas in neurofeedback training may provide the most potent means of enacting a change on visual perception and unconscious biases, potentially modulating the raw material of complex perception. Neural representations that give rise to prior expectations may be directly shifted in the direction of neurofeedback training, in the absence of an explicit expectation. This may in turn result in an enhancement of desirable responses during predictable as well as unpredictable tasks involving visual perception.

This type of perceptual shaping could be envisaged as being used for perceptual 'enhancements', where specific behaviour and perceptual routines could be preferentially trained. The potential benefit of using neurofeedback training would be to enable a broader repertoire of responses, based on modulating the level of activation within a predetermined area, which may have multiple, similar, overlapping neural representations. This was demonstrated by the study in Chapter 6, in which participants were preferentially shifted towards a specific *category* e.g. faces vs. houses, rather than being directed towards a specific face or house exemplar.

Clinically, neurofeedback training of higher-order visual brain regions could potentially be used in the treatment of thought disorders and phobias, to encourage perceptual biasing towards normative values, and away from intrusive and disruptive percepts. In a broader context, neurofeedback might be used in education and learning, to measure, track and accelerate skill acquisition. An example might be using neurofeedback training to specifically engage target brain regions, ahead of the use of linked educational materials utilising the trained brain regions. For a specific example, students might train visual brain regions prior to learning on a visual reasoning task. Further, knowledge of specific control regions implicated in particular types of learning e.g. SPL in neurofeedback and visual perceptual learning (81,540), could pave the way to specifically triggering activation in these regions in order to accelerate learning.

A tailored application of neurofeedback BCIs that is related to the above is the use of 'closed loop' training, where the user is not explicitly aware of the neurofeedback signal. De Bettencourt et al. (323) illustrated this system using rt-fMRI and category-specific attention. Brain activations related to either 'house' stimuli or 'face' stimuli were decoded as participants viewed superimposed stimuli. Participants were then guided towards attending to one or the other stimulus by dynamically altering the luminance of the stimuli in the visual display in relation to online decoding of brain activations in category specific ventral visual areas. This approach might potentially be extended to more complex cognitive processes with the aim of enhancing normal cognition or shifting a pathological brain state towards a normative one using an interactive display, an online assessment of correlated activation patterns (i.e. decoding a participant's mental state in real time) and continuous feedback in response. Using such an adaptive closed-loop feedback approach, the current mental state could be covertly guided towards a target state. A separate application outside the clinical arena might be during skill acquisition as alluded to above, where neural activation patterns could be monitored and guided towards 'decoded' neural states previously established as being associated with expertise (548).

An exploratory opportunity afforded by neurofeedback-guided shaping of neural representations would to prepare specific regions of cortex prior to 'fitting' a tailored neural prosthesis, which could be physical, cognitive or psychological. Preliminary work (549) indicates that participants can differentially and simultaneously modulate activation in left and right motor cortex as a result of neurofeedback-based training. The purpose of this work was to address the alteration in balance of hemispheric activity in the motor cortex after stroke.

On the basis on of this and related work (94), if motor cortex activity could be manipulated so specifically, it may then be possible to train activity in a precise region of the motor cortex prior to fitting a prosthesis. For example, neurofeedback-based training of the activation in the hand area motor region (i.e. hand 'knob' along the precentral sulcus (550)) could be performed prior to fitting a prosthetic hand. This would preferentially and pre-emptively engage necessary neural processes, which could then be further selectively trained and pruned through prosthetic hand use, and neurofeedback.

An analogous sensory prosthesis can also be envisaged, such as training activation in auditory brain regions prior to fitting a hearing aid, and/or audiology training. An alternative approach to a cognitive prosthesis based on closed-loop neurofeedback might be an enhanced human/computer interaction, in which through a 'co-adaptive' system (551), the BCI would learn from the environment, as well as from explicit input from the BCI user, resulting in continuous synergistic adaptation (552). A related concept is currently the focus of research by the Defense Advanced Research Projects Agency (DARPA) (548). For example, a computer based 'intelligent' threat detection system would engage programmed automated routines, as well as concurrently interacting with, and selectively training a human user to respond more rapidly to a particular threat through neurofeedback enhancement of specific cortical responses.

7.5.2 A 'hardware upgrade': 'Brain-in-the-loop' applications

A useful term coined in a recent publication in relation to BCIs, 'brain-in-the-loop' (548), highlights an important and potentially very powerful application of BCI technologies; rather than a behaviour being targeted for appraisal, modulation and shaping towards a desirable endpoint, brain function and activation are directly accessed and utilised (see Fig. 7-1).

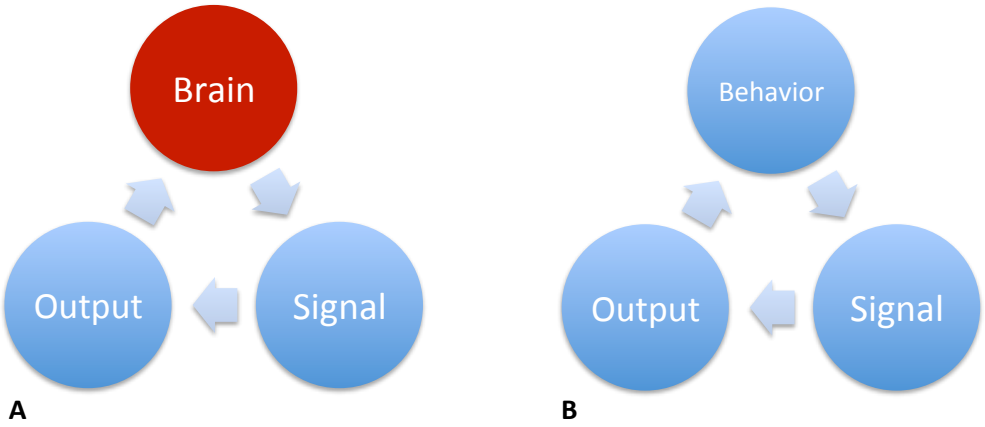

Fig 7-1. Illustration of a possible paradigm for BCI technology relying on the 'brain-in-the-loop' concept. As shown in (A), the brain is now the point of

access for the BCI, as compared to the standard paradigm (B), which instead accesses behaviour as the correlate of brain function.

I have discussed BCI-driven modulation of neural activation as the focal concept in the previous section. I will now consider 'hardware' approaches, that is, applications that may directly alter the structure of the brain, or may work alongside or 'extend' normal brain function.

7.5.3 Triggering adaptive plasticity to restore function

The term plasticity has been applied to a number of different measures of brain structure and function , and it is important to establish the metric associated with proposed or purported changes in plasticity for the relevant approach (e.g. neuroimaging or neurophysiological changes, and behavioural or structural changes), and the manner in which they interact. For example, by training a specific brain region or regions, changes in functional imaging measures have been correlated with changes in behaviour. Subramaniam et al. (13) demonstrated improvements on motor measures in PD patients following learned upregulation of SMA. Correlated alterations were noted in the levels of functional activation in the SMA, as well as other network brain regions implicated in PD, such as the subthalamic nucleus and the globus pallidus interna.

Shibata et al. (509) recorded changes in the function of the primary visual cortex, by training participants to produce brain activations matched to a particular grating stimulus orientation. Participants were then more likely to identify the same grating orientation, even when not pre-emptively evoking cortical activations associated with the training state, leading the authors to implicate changes in plasticity of V1 following perceptual learning. This finding with primary visual cortex is analogous to the findings in Chapter 6 that occurred instead in higher-order visual cortex. BR dynamics were altered following training, even when participants were not actively up-regulating the trained brain regions. These behavioural and correlated functional changes in brain activation have been considered direct examples of brain plasticity.

An important question that was addressed by the study in Chapter 6, was whether functional changes produce associated structural changes in grey and/or white matter, representing a novel finding in the literature.

Evidence for matched functional and structural plasticity has been reported from studies of motor learning in rodent models. Here, functional reorganisation of cortical motor maps, as well as structural changes that include grey matter glial changes and synaptogenesis, and increased myelination of relevant white matter pathways have been shown (553). In humans, cognitive training of healthy participants on working memory tasks was correlated with increased fractional anisotropy (FA) in white matter tracts adjacent to working memory regions, such as the body of the corpus callosum, and the IPS (554). FA has been linked with the degree of myelination, and is regarded as a marker of white matter integrity. Taken together these findings provide for an intriguing therapeutic possibility that relies on such structural changes, namely, could structural brain injury caused by disease or trauma be restored by focused neurofeedback training of brain activation in or near to the damaged brain region? Intuitively producing measurable structural brain changes through training would confirm a causal link between training and desirable changes in function and behaviour.

An alternative approach may be to target compensation for damaged brain function by training compensatory connectivity in associated network pathways. This may not necessarily only rely on overt structural changes, and in fact has been offered as an explanation by Subramaniam et al. (13), for the therapeutic benefit observed following neurofeedback modulation of the SMA in PD patients.

7.5.4 Decoding brain signals to enable function

Component cognitive functions are internal and ‘hidden’, and are typically experimentally assessed in an indirect manner. In order to ‘power’ a cognitive BCI, internal states comprising cognitive functions such as visual working memory, attention etc. must be accessed, a challenge well addressed by the use of fMRI,

currently the best means of identifying where in the brain, activation related to a specific cognitive process is occurring.

By employing classification techniques on brain activations measured with fMRI, the aim is to identify what brain function is occurring in relation to its cortical/subcortical location. This is in fact the first step in the design of any BCI – identification and deciphering of target brain activation (see Fig. 7-2A). These 'decoded' signals can then be used as the translated surrogate of the cognitive function, ideally providing a real-time view of brain mechanics occurring in its service. The optimal evolution of this technique would be to parse complex, hidden brain states into intuitive, semantic descriptors, in effect creating a brain 'translation' device. Given the current state-of-the-art, fMRI at 3T can only be used as a relatively coarse enabler of function, by decoding brain activation for use in the communication and control of a BCI. I achieved a reasonable degree of success with the online decoding of spatial attention from higher order visual areas (individual accuracies of 70%), following inclusion of multiple sources of information for decoding.

The emerging challenge then is achieving sufficient resolution with the brain imaging technique, to accurately extract all of the neural activity related to a specific cognitive function, and employing techniques that identify and reconstruct the neural activity. The ideal decoding outcome for the experiment in Chapter 6 might be hypothesised as being able to determine in real time what specific object exemplar was being attended to, in which quadrant, on an exemplar-by-exemplar basis. The required rates for accurate decoding with a BCI to be operational for communication have been suggested to be 70% (467–469). This was initially suggested in the context of BCI using neurophysiological measures. Higher rates of online decoding for spatial attention have been achieved with rt-fMRI at 7T (10,19,43,311), although this was obtained with primary visual cortex, using simple high contrast stimuli e.g. chequered gratings in two directions and high luminescence arrows in four directions.

More recently, attempts have been made to transform voxel-based information from individual brain activity into a common space, using statistical learning methods to subsequently provide flexible decoding of visual images across, rather than just within, individuals (555). The neural decoder was able to successfully predict fMRI

activity patterns of a target participant for unseen visual stimuli (structured or random patches), having being provided with fMRI activity patterns from a different participant. This work demonstrated the first steps towards translating a common 'neural code' across individuals (555). Although this work was demonstrated with primary visual cortex, in keeping with the flexibility of the machine learning methodology it is dependent on, it could be applicable to higher-order brain regions. For example, it could be applied to the decoding of more complex functions such as social behaviours and even aspects of personality. Such an application could provide insights into individual differences in brain function on the one hand, and provide means of crossing boundaries between individuals presented by language, and distance. Direct brain-to-brain communication (556–558) utilising a common neural code and foregoing the requirement for overt communication, is already an active area of research, and would represent an ultimate destination for brain decoding and the enabling of function.

7.5.5 Combining invasive and non-invasive BCI technology

The work in this thesis is specifically concerned with use of fMRI as means of developing a non-invasive cognitive BCI. There are a number of advantages to the use of non-invasive technology in addition to avoiding the need for a surgical procedure e.g. ease of replacement, upkeep and modification of the BCI. On the other hand an ideal BCI device needs to be small and discrete, with a minimal aesthetic profile (559). Current non-invasive technologies, taking into account the most portable examples on the market, such as wireless, dry electrode EEG headsets (560–565), do not easily fit this criteria, and remain relatively cumbersome. Implantable technologies have been investigated with the aim of providing a fully concealed, portable, wireless device, and have been shown to be preferable from the point of view of the patient or participant (559). A recent documented success has been with the 'BrainGate' BCI (v2.0), using a microelectrode array implanted 5 years previously in a patient with incomplete locked-in-syndrome. The array, implanted in the hand/arm area of the motor cortex, enables control via a point-and-click interface with a novel hexadecimal-based radial keyboard (536). A persistent

issue with invasive BCI technologies, given the intrinsic aim of maintaining long-term implantation, is, however, hardware failure and implant infection (566,567). These, and other issues related to interface failure vs. reliability are currently being addressed by the DARPA-funded Reliable Neural Interface Technology (RE-NET) program involving the joint effort of a number of top-tier US university research institutions (548).

A useful approach to these specific issues related to BCI's is to combine the strengths of non-invasive and implantable techniques. A potential pipeline might be to first use decoding algorithms to identify the most appropriate regions of cortex that are functionally relevant for an assigned task. Online neurofeedback training with or without a closed loop system could then be used to optimise the neural representations in relation to the functional requirement. In the third step, an implantable BCI would be inserted in a focal or distributed fashion, in anatomical alignment with the non-invasively identified and trained cortical ROIs. Finally a fourth step could be envisaged where adaptive neurofeedback via the implanted BCI could be used to train and maintain specific neuronal populations or activations, depending on the resolution of the implanted BCI. A simple proof-of-principle has been demonstrated at 7T using online decoding of spatial attention in primary visual cortex prior to implantation of a subdural ECoG electrode grid array (10).

Non-invasive BCI's may provide an alternative means of stimulating brain activation, in a manner analogous to more direct surgical approaches e.g. deep brain stimulation (DBS) with implantable electrodes (568). DBS has been successfully used to stimulate a number of subcortical targets for therapeutic effects on behaviour, cognition, autonomic physiology and movement (for reviews see 65, 66). Conceptually, rt-fMRI neurofeedback might provide a means of access to subcortical targets in a similar fashion. Indeed recent work has shown successful online modulation of the ventral tegmental area (11), a region which is currently being used as a surgical DBS target for the treatment of cluster headache and refractory depression (571). Further, non-invasive modulation of cortical regions have been shown to produce downstream effects on subcortical regions, with associated behavioural changes in PD (13). It may be possible to utilise rt-fMRI 'non-invasive brain stimulation' method to treat some of the perceptual abnormalities in PD,

incorporating the benefits of mental imagery training linked to rt-fMRI mediated access of cortical and subcortical ROIs e.g. ventral visual areas (572) or the Nucleus Basalis of Meynert (573).

A example pipeline for rt-fMRI 'non-invasive brain stimulation' method might be to first identify PD patients who have been selected for DBS surgery, prior to surgical implantation of the deep brain electrodes, and to then use rt-fMRI neurofeedback training to modulate specific cortical and subcortical nodes, as well as networks joining the two, e.g. SMA, STN, and the hyperdirect pathway (13,574–576). The aim would be to potentiate specific pathways and targets prior to surgery, potentially providing an additive effect to surgery, as well as optimising surgical targets. This approach might eventually provide a standalone means of treatment, which would be suitable for patients not deemed fit for surgery, or those opting not to undertake the attendant risks of a surgical procedure. Following surgical implantation of the electrodes, it is possible to 'externalise' the leads for a short period of time, and undertake local field potential recordings from the implanted STN. This would provide an opportunity to examine the down stream effects of learned neurofeedback regulation of cortical targets i.e. SMA in relation to the surgical target i.e. STN. Indeed, this investigational pathway has formed the basis of a successful grant application to trial the use of neurofeedback training of SMA modulation in pre-operative PD patients, prior to STN-DBS surgery.

7.6 Neurosurgery and the BCI method – a therapeutic opportunity

The combination of the neurosurgical and neuroimaging approaches provides fertile ground for the development and application of neurofeedback and BCI ideas and technologies. Most frequently, implanted stimulation and recording devices such as DBS electrodes, and ECoG with subdural electrodes arrays, allow for direct measurements of brain function from the single neuron level up to the mesoscopic scale, sampling focal populations of neuronal activity (e.g. populations of 10^5 neurones are measured with ECoG). This can be used to enable BCI effector

control, as well as providing a unique opportunity for measuring brain function in response to invasive and non-invasive stimulation modulation.

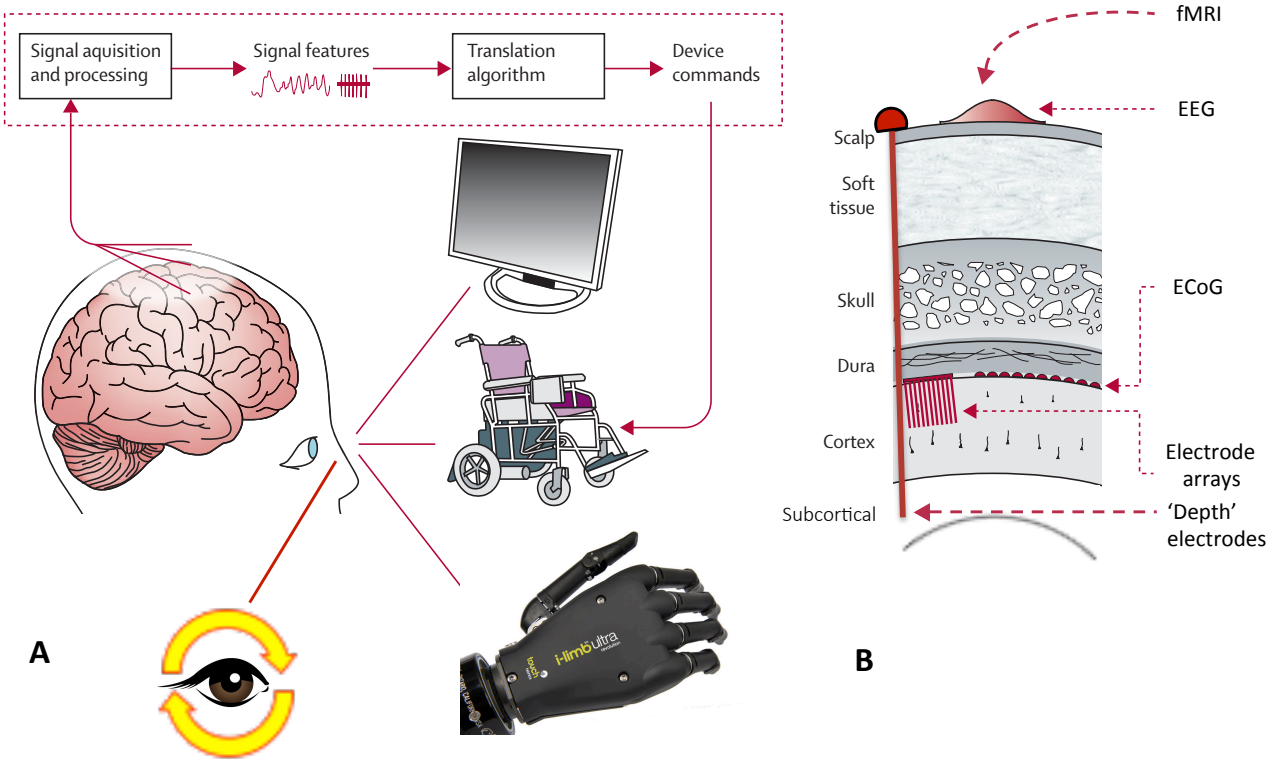

Fig 7-2A. Schematic of the components parts of a typical BCI set-up.

This includes signal acquisition from the brain, and translation into effector outputs which may be control of a computer or prosthetics such as a wheelchair or robotic hand, and for self-modulation of perceptual processes (Adapted from Wolpaw[90]).

Fig 7-2B. Anatomical layout illustrating the multiple points of access for BCI interventions from the scalp to the deep subcortical regions.

In the preceding section, I have described one potential example (i.e. rt-fMRI 'non-invasive brain stimulation' neurofeedback training in pre-op STN-DBS) of the manner in which routine neurosurgical implantation of prosthetics can be leveraged as an opportunity for examining the modulation of brain activation in relation to neurofeedback training and therapeutic behavioural effects. DBS is a versatile

surgical technology that can be used in a number of ways in relation to BCI technologies and neurofeedback. For example, adaptive closed-loop stimulation with DBS is an active area of research (577,578), which may enable tailored and more efficient stimulation paradigms, improving therapeutic benefit as well increasing hardware lifetime in-vivo. Online measurements of an implanted closed loop system using wireless measuring devices are currently in development; longitudinal in-vivo measurements of implanted brain regions could then be performed, potentially enabling on-going neurofeedback sessions for patients.

Recording from single neurons with depth electrodes implanted for seizure location identification in epilepsy patients have been used to investigate perceptual processes, such as responses in the temporal lobe when patients were shown celebrity face images (579), and lateral temporal cortex responses to speech sounds (580). ECoG arrays, also inserted at the time of surgery for seizure detection in epilepsy, have been used to provide proof-of-principle for computer-cursor control using auditory cortex activations (581), covert shifts of spatial attention in sensorimotor areas (432,476), and the decoding of speech production-related cortical activity (580,582). Most recently, an ECoG array was inserted de novo in a tetraplegic patient over sensorimotor cortex, to enable 3-dimensional BCI control of a computer cursor (583). Such studies are illustrative of the important insights that can be gained from the integration of neurosurgical and neuroimaging disciplines, with the point of overlap provided by the implantation of prosthetic devices that can directly stimulate or record from the brain.

The neurosurgical insertion of ICP bolts to provide continuous intracranial pressure measurements could also be employed as a novel opportunity to apply the BCI/ biofeedback concept for therapeutic benefit. Intracranial recordings of autonomic physiology may be used as a target for neurofeedback-guided therapies. A significant conclusion from the completion of the work in Chapter 6 was the potential power of the neurofeedback technique as a therapeutic tool. Neurofeedback can be viewed as a variant of biofeedback. As established in the introductory and method chapters, the biofeedback approach was first demonstrated in the 1960s, with voluntary bi-directional control of heart rate via an analogue visual display which provided the biological signal to be controlled (584). The fundamental conclusion of

this and subsequent work using measures of internal autonomic activity, including cerebral physiology e.g. control of regional cerebral blood flow and end tidal CO_2 to treat migraine and epilepsy (585–588), was that biofeedback enables voluntary control of autonomic physiology. This may be mediated by a putative central autonomic network, and its interaction with the peripheral autonomic nervous system (589–591). It is possible to speculate that neurofeedback modulation of the BOLD signal (i.e. a haemodynamic measure), through the use of mental imagery, may be based upon control of cerebral vasculature. On the basis of these insights, and those developed in this thesis in relation to the neurofeedback concept, further work (see Fig. 7-3 and 7-4) and a scientific proposal have been submitted (592). The proposal seeks to demonstrate that biofeedback-based control of intracranial pressure (ICP) is possible, and may be used for patients with ICP-related pathology e.g. intracranial idiopathic hypertension. The scientific justification is based on the assumption that if biofeedback control of cardiovascular and cerebral vasculature is possible, other dependent measures of cerebral physiology, such as ICP, may also be amenable to volitional control.

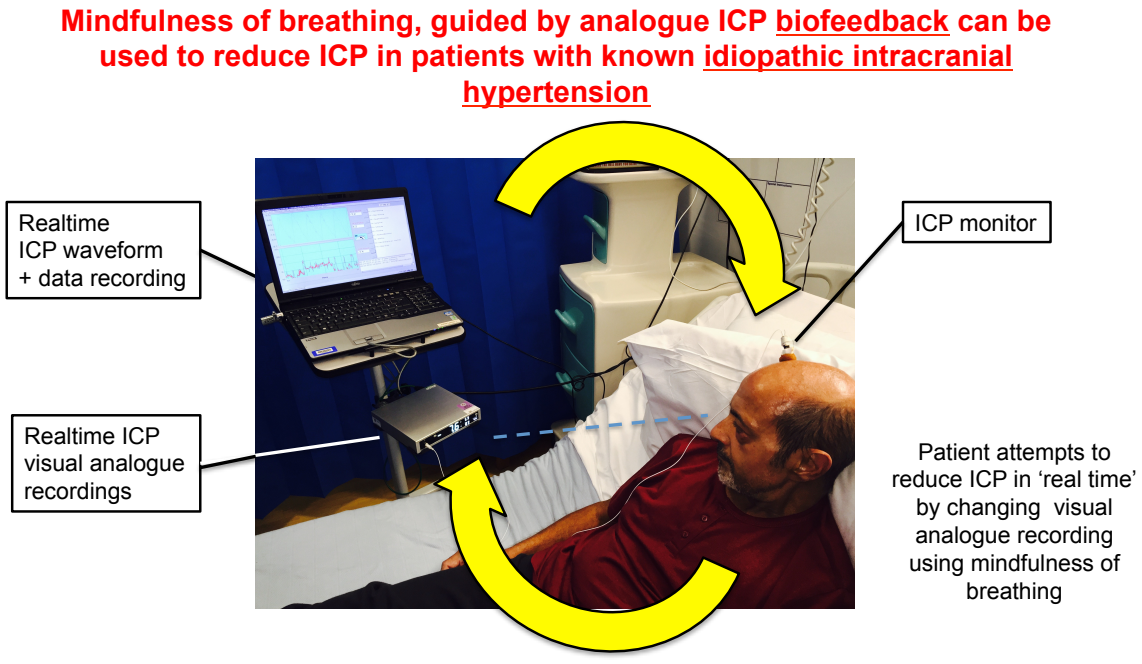

Fig 7-3. ICP Biofeedback set-up. Patient shown with ICP bolt inserted and connected to a computer display showing ICP wave form, and an analogue display of absolute ICP values in real-time. The patient attempts to reduce the number of the ICP value using biofeedback control.

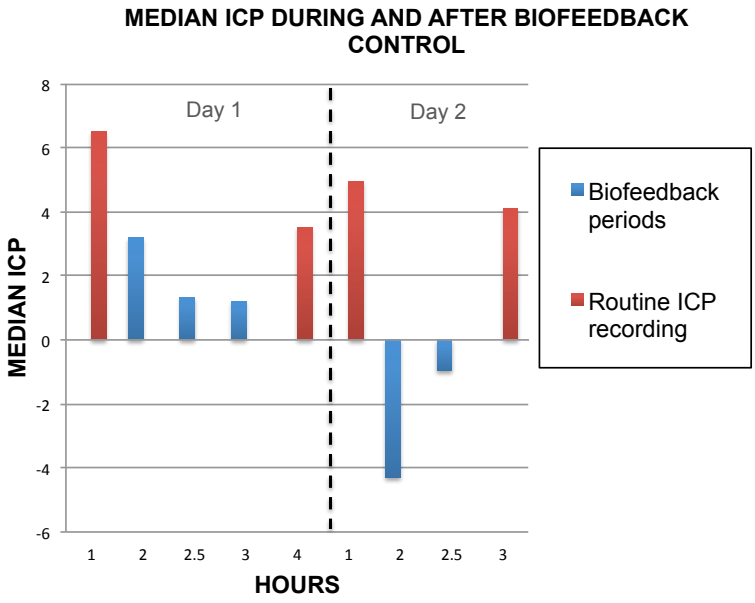

Fig 7-4. Preliminary data for ICP biofeedback study. Data from a single IIH patient, illustrating ICP changes during concurrent ICP biofeedback and mindfulness of breathing. On the x-axis are the number of hours over which ICP was monitored, and the y-axis shows median ICP values (mmHg). These preliminary results indicate that the patient learned to voluntarily reduce their ICP values using biofeedback control.

7.7 Summary and final conclusions

In this thesis, I presented empirical evidence for the use of higher-order visual areas, and their functioning in visual cognition and perception, as potential targets for non-invasive BCI using rt-fMRI. In the first half of my thesis I demonstrated that it is possible to reliably decode the direction of covertly deployed spatial attention, online and offline, exploiting the use of spatial, and object category information, in addition to temporal information of visual stimuli arising from the novel use of m-sequences. In the second half of my thesis, after successfully establishing an rt-fMRI analysis pathway, I demonstrated that category specific visual perception can be biased by neurofeedback training of ventral visual areas. Perceptual dynamics during a BR

paradigm were altered, providing a novel conceptualisation of a long-standing observation regarding BR, known as Levelt's 2nd proposition.

The use of category specific visual areas in this thesis confirmed that they can provide accurate and reliable decoding, may be manipulated across a range of cognitive functions, and offer an access point to complex functions such as perception. I further investigated the underlying mechanics of neurofeedback, utilising a number of cutting-edge analysis techniques, and implicated the involvement of top-down control by a putative modulatory region, SPL. This finding may provide empirical evidence for an extant theory of biofeedback, known as 'Lacroix's dual step process'.

From an operational point of view, I have discussed the potential use of a rt-fMRI set-up with neurofeedback as a non-invasive BCI, in health and disease. I also covered biofeedback, which relies on concurrent electronic measurement to 'close' the loop- to this end it is by definition a BCI. Nonetheless, one of long-term aims of biofeedback/ neurofeedback learning is to automatise learned control of the trained measure, and achieve control of the target physiology in the absence of concurrent feedback.

If function can be restored and maintained in the absence of prosthetic implants, this would clearly be optimal, but may not always be possible. To this end, biofeedback, decoding and neurofeedback, and non-invasive and implantable BCIs form a continuum, which provide a number of flexible options for enabling and restoring brain function, which is summarised in Fig 7-2B.

To conclude, the research in the areas covered by the work in this thesis is at an exciting stage. Technological advances in computing, imaging and engineering are being more readily combined with human biology, giving rise to the emerging yet rapidly developing field of BCI, with the principal aim of enhancing brain function in health and restoring brain function in disease.

8. REFERENCES

1. Lebedev M a, Nicolelis M a L. Brain-machine interfaces: past, present and future. Trends Neurosci. 2006 Sep;29(9):536–46.
2. Sitaram R, Caria A, Veit R, Gaber T, Rota G, Kuebler A, et al. FMRI brain-computer interface: a tool for neuroscientific research and treatment. Comput Intell Neurosci. 2007 Jan;2007:25487.
3. Daly JJ, Wolpaw JR. Brain-computer interfaces in neurological rehabilitation. Lancet Neurol. Elsevier Ltd; 2008 Nov;7(11):1032–43.
4. Birbaumer N. Breaking the silence: brain-computer interfaces (BCI) for communication and motor control. Psychophysiology. 2006 Nov;43(6):517–32.
5. Tehovnik EJ, Woods LC, Slocum WM. Transfer of information by BMI. Neuroscience. IBRO; 2013 Dec 26;255:134–46.
6. Min B-K, Marzelli MJ, Yoo S-S. Neuroimaging-based approaches in the brain-computer interface. Trends Biotechnol. Elsevier Ltd; 2010 Nov;28(11):552–60.
7. Lee J-H, Ryu J, Jolesz F a, Cho Z-H, Yoo S-S. Brain-machine interface via real-time fMRI: preliminary study on thought-controlled robotic arm. Neurosci Lett. 2009 Jan 23;450(1):1–6.
8. Stanslaski S, Cong P, Carlson D, Santa W, Jensen R, Molnar G, et al. An implantable bi-directional brain-machine interface system for chronic neuroprosthesis research. Conf Proc IEEE Eng Med Biol Soc. 2009 Jan;2009:5494–7.
9. Tankus A, Fried I, Shoham S. Cognitive-motor brain-machine interfaces. J Physiol Paris. 2014 Feb;108(1):38–44.
10. Andersson P, Pluim JPW, Siero JCW, Klein S, Viergever M a, Ramsey NF. Real-time decoding of brain responses to visuospatial attention using 7T fMRI. PLoS One. 2011 Jan;6(11):e27638.
11. Sulzer J, Sitaram R, Blefari ML, Kollias S, Birbaumer N, Stephan KE, et al. Neurofeedback-mediated self-regulation of the dopaminergic midbrain. Neuroimage. 2013;83:817–25.

12. Koush Y, Rosa MJ, Robineau F, Heinen K, W Rieger S, Weiskopf N, et al. Connectivity-based neurofeedback: dynamic causal modeling for real-time fMRI. Neuroimage. Elsevier Inc.; 2013 Nov 1;81:422–30.

13. Subramanian L, Hindle J V, Johnston S, Roberts M V, Husain M, Goebel R, et al. Real-time functional magnetic resonance imaging neurofeedback for treatment of Parkinson's disease. J Neurosci. 2011 Nov 9;31(45):16309–17.

14. Monti MM, Vanhaudenhuyse A, Coleman MR, Boly M, Pickard JD, Tshibanda L, et al. Willful modulation of brain activity in disorders of consciousness. N Engl J Med. 2010 Feb 18;362(7):579–89.

15. Astrand E, Wardak C, Ben Hamed S. Selective visual attention to drive cognitive brainâ€“machine interfaces: from concepts to neurofeedback and rehabilitation applications. Front Syst Neurosci. 2014 Aug 12;8(August):1–16.

16. Scharnowski F, Hutton C, Josephs O, Weiskopf N, Rees G. Improving visual perception through neurofeedback. J Neurosci. 2012 Dec 5;32(49):17830–41.

17. Shibata K, Watanabe T, Sasaki Y, Kawato M. Perceptual learning incepted by decoded fMRI neurofeedback without stimulus presentation. Science. 2011 Dec 9;334(6061):1413–5.

18. Geranmayeh F, Brownsett SLE, Wise RJS. Task-induced brain activity in aphasic stroke patients: what is driving recovery? Brain. 2014 Oct;137(Pt 10):2632–48.

19. Andersson P, Ramsey NF, Viergever M a, Pluim JPW. 7T fMRI reveals feasibility of covert visual attention-based brain-computer interfacing with signals obtained solely from cortical grey matter accessible by subdural surface electrodes. Clin Neurophysiol. International Federation of Clinical Neurophysiology; 2013 Nov;124(11):2191–7.

20. Mushahwar VK, Hochberg LR, Donoghue JP. BCI Meeting 2005 — Workshop on Clinical Issues and. 2006;14(2):131–4.

21. Hochberg LR, Serruya MD, Friehs GM, Mukand J a, Saleh M, Caplan AH, et al. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature. 2006 Jul 13;442(7099):164–71.

22. Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, et al.

Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. Nature. Nature Publishing Group; 2012 May 17;485(7398):372–5.

23. Brumberg JS, Nieto-Castanon A, Kennedy PR, Guenther FH. Brain-Computer Interfaces for Speech Communication. Speech Commun. 2010 Apr 1;52(4):367–79.

24. Chaudhary U, Birbaumer N, Curado MR. Brain-Machine Interface (BMI) in paralysis. Ann Phys Rehabil Med. Elsevier Masson SAS; 2015;58(1):9–13.

25. McCane LM, Heckman SM, McFarland DJ, Townsend G, Mak JN, Sellers EW, et al. P300-based brain-computer interface (BCI) event-related potentials (ERPs): People with amyotrophic lateral sclerosis (ALS) vs. age-matched controls. Clin Neurophysiol. International Federation of Clinical Neurophysiology; 2015;126(11):2124–31.

26. Wolpaw J. Brain-computer interfaces : principles and practice. Oxford ;;New York: Oxford University Press; 2012.

27. Sellers EW, Ryan DB, Hauser CK. Noninvasive brain-computer interface enables communication after brainstem stroke. Sci Transl Med. 2014;6(257):257re7.

28. Ogawa S, Lee TM, Kay AR, Tank DW. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc Natl Acad Sci. 1990 Dec 1;87(24):9868–72.

29. Ogawa S, Lee TM, Nayak AS, Glynn P. Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. Magn Reson Med. 1990 Apr;14(1):68–78.

30. Seriès P, Seitz AR. Learning what to expect (in visual perception). Front Hum Neurosci. 2013 Jan;7(October):668.

31. Cox RW, Jesmanowicz A, Hyde JS. Real-Time Functional Magnetic Resonance Imaging. Magn Reson Med. 1995 Feb;33(2):230–6.

32. Cox RW, Jesmanowicz A. Real-time 3D image registration for functional MRI. Magn Reson Med. 1999;42(6):1014–8.

33. Mathiak K, Posse S. Evaluation of motion and realignment for functional magnetic resonance imaging in real time. Magn Reson Med. 2001

Jan;45(1):167–71.

34. Smyser C, Grabowski TJ, Frank RJ, Haller JW, Bolinger L. Real-time multiple linear regression for fMRI supported by time-aware acquisition and processing. Magn Reson Med. 2001 Feb;45(2):289–98.

35. Goebel R. Cortex-based real-time fMRI. Neuroimage. 2001;13(6):S129–S129).

36. Goddard N, Hood G, Cohen J, Eddy W, Genovese C, Noll D, et al. Online Analysis of Functional MRI Datasets on Parallel Platforms. J Supercomput. Kluwer Academic Publishers; 11(3):295–318.

37. Voyvodic JT. Real-time fMRI paradigm control, physiology, and behavior combined with near real-time statistical analysis. Neuroimage. 1999 Aug;10(2):91–106.

38. Haynes J-D, Rees G. Decoding mental states from brain activity in humans. Nat Rev Neurosci. 2006 Jul;7(7):523–34.

39. Sorger B, Reithler J, Dahmen B, Goebel R. A real-time fMRI-based spelling device immediately enabling robust motor-independent communication. Curr Biol. Elsevier Ltd; 2012 Jul 24;22(14):1333–8.

40. Yoo S-S, Fairneny T, Chen N-K, Choo S-E, Panych LP, Park H, et al. Brain–computer interface using fMRI: spatial navigation by thoughts. Neuroreport. 2004 Jul;15(10):1591–5.

41. Saygin AP, Sereno MI. Retinotopy and attention in human occipital, temporal, parietal, and frontal cortex. Cereb Cortex. 2008 Sep;18(9):2158–68.

42. Bahramisharif A, van Gerven M, Heskes T, Jensen O. Covert attention allows for continuous control of brain-computer interfaces. Eur J Neurosci. 2010 Apr;31(8):1501–8.

43. Andersson P, Ramsey NF, Raemaekers M, Viergever M a, Pluim JPW. Real-time decoding of the direction of covert visuospatial attention. J Neural Eng. 2012 Aug;9(4):045004.

44. Andersson P, Ramsey NF, Pluim JPW, Viergever M a. BCI control using 4 direction spatial visual attention and real-time fMRI at 7T. Conf Proc IEEE Eng Med Biol Soc. 2010 Jan;2010:4221–5.

45. Andersson P, Pluim JPW, Viergever M a, Ramsey NF. Navigation of a

telepresence robot via covert visuospatial attention and real-time fMRI. Brain Topogr. 2013 Jan;26(1):177–85.

46. Zenon A, Filali N, Duhamel J-R, Olivier E. Salience representation in the parietal and frontal cortex. J Cogn Neurosci. 2010 May;22(5):918–30.

47. Gottlieb J. From thought to action: the parietal cortex as a bridge between perception, action, and cognition. Neuron. 2007 Jan 4;53(1):9–16.

48. Cichy RM, Chen Y, Haynes J-D. Encoding the identity and location of objects in human LOC. Neuroimage. Elsevier Inc.; 2011 Mar 1;54(3):2297–307.

49. Golomb JD, Kanwisher N. Higher level visual cortex represents retinotopic, not spatiotopic, object location. Cereb Cortex. 2012 Dec;22(12):2794–810.

50. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, et al. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. 2006;442(July).

51. Lang PJ, Twentyman CT. Learning to control heart rate: effects of varying incentive and criterion of success on task performance. Psychophysiology. 1976;13(5):378–85.

52. Fetz EE. Operant Conditioning of Cortical Unit Activity. Science (80-). 1968;163(3870):955–8.

53. Jackson A, Mavoori J, Fetz EE. Long-term motor cortex plasticity induced by an electronic neural implant. Nature. 2006 Nov 2;444(7115):56–60.

54. Wyrwicka W, Sterman MB. Instrumental conditioning of sensorimotor cortex EEG spindles in the waking cat. Physiol Behav. 1968 Sep;3(5):703–7.

55. Kamiya J, Callaway E, Yeager CL. Visual evoked responses in subjects trained to control alpha rhythms. Psychophysiology. 1969. p. 683–95.

56. Kamiya J. The First Communications About Operant Conditioning of the EEG. J Neurother. 2011 Feb 22;15(1):65–73.

57. Rota G, Sitaram R, Veit R, Erb M, Weiskopf N, Dogil G, et al. Self-regulation of regional cortical activity using real-time fMRI: the right inferior frontal gyrus and linguistic processing. Hum Brain Mapp. 2009 May;30(5):1605–14.

58. deCharms RC, Christoff K, Glover GH, Pauly JM, Whitfield S, Gabrieli JD.

Learned regulation of spatially localized brain activation using real-time fMRI. Neuroimage. 2004 Jan;21(1):436–43.

59. deCharms RC, Maeda F, Glover GH, Ludlow D, Pauly JM, Soneji D, et al. Control over brain activation and pain learned by using real-time functional MRI. Proc Natl Acad Sci U S A. 2005 Dec 20;102(51):18626–31.

60. Rance M, Ruttorf M, Nees F, Schad LR, Flor H. Real time fMRI feedback of the anterior cingulate and posterior insular cortex in the processing of pain. Hum Brain Mapp. 2014 Dec;35(12):5784–98.

61. Emmert K, Breimhorst M, Bauermann T, Birklein F, Van De Ville D, Haller S. Comparison of anterior cingulate vs. insular cortex as targets for real-time fMRI regulation during pain stimulation. Front Behav Neurosci. 2014 Jan;8:350.

62. Scharnowski F, Veit R, Zopf R, Studer P, Bock S, Göbel R, et al. Manipulating motor performance and memory through neurofeedback. :1–34.

63. Zotev V, Krueger F, Phillips R, Alvarez RP, Simmons WK, Bellgowan P, et al. Self-regulation of amygdala activation using real-time FMRI neurofeedback. PLoS One. 2011 Jan 8;6(9):e24522.

64. Hamilton JP, Glover GH, Hsu J-J, Johnson RF, Gotlib IH. Modulation of subgenual anterior cingulate cortex activity with real-time neurofeedback. Hum Brain Mapp. 2011 Jan;32(1):22–31.

65. Yoo S-S, Jolesz F a. Functional MRI for neurofeedback: feasibility study on a hand motor task. Neuroreport. 2002 Aug 7;13(11):1377–81.

66. Haller S, Kopel R, Jhooti P, Haas T, Scharnowski F, Lovblad K-O, et al. Dynamic reconfiguration of human brain functional networks through neurofeedback. Neuroimage. 2013 Nov 1;81:243–52.

67. Haller S, Birbaumer N, Veit R. Real-time fMRI feedback training may improve chronic tinnitus. Eur Radiol. 2010 Mar;20(3):696–703.

68. Sitaram R, Caria A, Veit R, Gaber T, Ruiz S, Birbaumer N. Volitional control of the anterior insula in criminal psychopaths using real-time fMRI neurofeedback: a pilot study. Front Behav Neurosci. 2014;8(344).

69. Linden DEJ, Habes I, Johnston SJ, Linden S, Tatineni R, Subramanian L, et al. Real-time self-regulation of emotion networks in patients with depression.

PLoS One. 2012 Jan;7(6):e38115.

70. Chang M, Iizuka H, Naruse Y, Ando H, Maeda T. Unconscious learning of auditory discrimination using mismatch negativity (MMN) neurofeedback. Sci Rep. 2014 Jan;4:6729.

71. Brownsett SLE, Warren JE, Geranmayeh F, Woodhead Z, Leech R, Wise RJS. Cognitive control and its impact on recovery from aphasic stroke. Brain. 2014 Jan;137(Pt 1):242–54.

72. Strehl U. What learning theories can teach us in designing neurofeedback treatments. Front Hum Neurosci. 2014 Jan;8(November):894.

73. Power JD, Petersen SE. Control-related systems in the human brain. Curr Opin Neurobiol. Elsevier Ltd; 2013 Apr;23(2):223–8.

74. Ruiz S, Rana M, Sass K. Brain network connectivity and behaviour enhancement: a fMRI-BCI study. 17th Annu Meet …. 2011;

75. Ruiz S, Lee S, Soekadar SR, Caria A, Veit R, Kircher T, et al. Acquired self-control of insula cortex modulates emotion recognition and brain network connectivity in schizophrenia. Hum Brain Mapp. 2013 Jan;34(1):200–12.

76. Ruiz S, Buyukturkoglu K, Rana M, Birbaumer N, Sitaram R. Real-time fMRI brain computer interfaces: Self-regulation of single brain regions to networks. Biol Psychol. Elsevier B.V.; 2014 Jan;95:4–20.

77. Enriquez-Geppert S, Huster RJ, Scharfenort R, Mokom ZN, Vosskuhl J, Figge C, et al. The morphology of midcingulate cortex predicts frontal-midline theta neurofeedback success. Front Hum Neurosci. 2013 Jan;7:453.

78. Halder S, Varkuti B, Bogdan M, Kübler A, Rosenstiel W, Sitaram R, et al. Prediction of brain-computer interface aptitude from individual brain structure. Front Hum Neurosci. 2013 Jan;7:105.

79. Loui P, Li HC, Schlaug G. White matter integrity in right hemisphere predicts pitch-related grammar learning. Neuroimage. 2011 Mar 15;55(2):500–7.

80. Basak C, Voss MW, Erickson KI, Boot WR, Kramer AF. Regional differences in brain volume predict the acquisition of skill in a complex real-time strategy videogame. Brain Cogn. 2011 Aug;76(3):407–14.

81. Scharnowski F, Rosa MJ, Golestani N, Hutton C, Josephs O, Weiskopf N, et

al. Connectivity changes underlying neurofeedback training of visual cortex activity. PLoS One. 2014 Jan;9(3):e91090.

82. Emmert K, Kopel R, Sulzer J, Br??hl AB, Berman BD, Linden DEJ, et al. Meta-analysis of real-time fMRI neurofeedback studies using individual participant data: How is brain regulation mediated? Neuroimage. Elsevier Inc.; 2016;124:806–12.

83. Menon V. Salience Network. Brain Mapping: An Encyclopedic Reference. Elsevier Inc.; 2015. 597-611 p.

84. Skinner BF. The operational analysis of psychological terms.

85. Weiskopf N. Real-time fMRI and its application to neurofeedback. Neuroimage. Elsevier Inc.; 2012 Aug 15;62(2):682–92.

86. Birbaumer N, Ruiz S, Sitaram R. Learned regulation of brain metabolism. Trends Cogn Sci. Elsevier Ltd; 2013 May 7;17(6):295–302.

87. Mathiak K a, Alawi EM, Koush Y, Dyck M, Cordes JS, Gaber TJ, et al. Social reward improves the voluntary control over localized brain activity in fMRI-based neurofeedback training. Front Behav Neurosci. 2015 Jan;9:136.

88. Bray S, Shimojo S, Doherty JPO. Direct Instrumental Conditioning of Neural Activity Using Functional Magnetic Resonance Imaging-Derived Reward Feedback. 2007;27(28):7498–507.

89. Mathiak KA, Koush Y, Dyck M, Gaber TJ, Alawi E, Zepf FD, et al. Social reinforcement can regulate localized brain activity. Eur Arch Psychiatry Clin Neurosci. 2010 Nov;260 Suppl :S132–6.

90. Lacroix JM. Mechanisms of biofeedback control: on the importance of verbal (conscious) processing. Conscious Self-Regulation. 1986;4:137–62.

91. Kober SE, Witte M, Ninaus M, Neuper C, Wood G. Learning to modulate one's own brain activity: the effect of spontaneous mental strategies. Front Hum Neurosci. 2013 Jan;7(October):695.

92. Garrison K a, Scheinost D, Worhunsky PD, Elwafi HM, Thornhill T a, Thompson E, et al. Real-time fMRI links subjective experience with brain activity during focused attention. Neuroimage. 2013 Nov 1;81:110–8.

93. Kropotov JD. Quantitative EEG, Event-Related Potentials and Neurotherapy.

Quantitative EEG, Event-Related Potentials and Neurotherapy. Elsevier; 2009. 161-180 p.

94. Koralek AC, Jin X, Long JD, Costa RM, Carmena JM. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. Nature. Nature Publishing Group; 2012 Mar 15;483(7389):331–5.

95. Palmer S. Vision science photons to phenomenology. Cambridge, Mass. : MIT Press,; 1999.

96. Grill-Spector K, Malach R. The human visual cortex. Annu Rev Neurosci. 2004 Jan;27:649–77.

97. Rossion B, Hanseeuw B, Dricot L. Defining face perception areas in the human brain: a large-scale factorial fMRI face localizer analysis. Brain Cogn. Elsevier Inc.; 2012 Jul;79(2):138–57.

98. Schwarzlose RF, Swisher JD, Dang S, Kanwisher N. The distribution of category and location information across object-selective regions in human visual cortex. Proc Natl Acad Sci U S A. 2008 Mar 18;105(11):4447–52.

99. Kanwisher N. Functional specificity in the human brain : A window into the functional architecture of the mind. 2010;

100. Gilbert CD, Sigman M. Brain states: top-down influences in sensory processing. Neuron. 2007 Jun 7;54(5):677–96.

101. Baluch F, Itti L. Mechanisms of top-down attention. Trends Neurosci. 2011 Apr;34(4):210–24.

102. Mishkin M, Ungerleider LG, Macko KA. Object vision and spatial vision: two cortical pathways. Trends Neurosci. 1983 Jan;6:414–7.

103. Goodale MA, Milner AD. Separate visual pathways for perception and action. Trends Neurosci. 1992 Jan;15(1):20–5.

104. Weiner KS, Grill-Spector K. Neural representations of faces and limbs neighbor in human high-level visual cortex: evidence for a new organization principle. Psychol Res. 2013 Jan;77(1):74–97.

105. Corbetta M, Shulman GL. Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci. 2002 Mar;3(3):201–15.

106. Hales JB, Brewer JB. Parietal and frontal contributions to episodic encoding of location. Behav Brain Res. 2013 Apr 15;243:16–20.

107. Qiu FT, Sugihara T, von der Heydt R. Figure-ground mechanisms provide structure for selective attention. Nat Neurosci. 2007;10(11):1492–9.

108. Fang F, Boyaci H, Kersten D. Border ownership selectivity in human early visual cortex and its modulation by attention. J Neurosci. 2009 Jan 14;29(2):460–5.

109. Zhang NR, von der Heydt R. Analysis of the context integration mechanisms underlying figure-ground organization in the visual cortex. J Neurosci. 2010 May 12;30(19):6482–96.

110. Zaretskaya N, Anstis S, Bartels A. Parietal cortex mediates conscious perception of illusory gestalt. J Neurosci. 2013 Jan 9;33(2):523–31.

111. Davis G, Driver J. Parallel detection of Kanizsa subjective figures in the human visual system. Nature. 1994;371(6500):791–3.

112. Larsson J, Heeger DJ. Two retinotopic visual areas in human lateral occipital cortex. J Neurosci. 2006 Dec 20;26(51):13128–42.

113. Cichy RM, Heinzle J, Haynes J-D. Imagery and perception share cortical representations of content and location. Cereb Cortex. 2012 Feb;22(2):372–80.

114. Karten A, Pantazatos SP, Khalil D, Zhang X, Hirsch J. Dynamic coupling between the lateral occipital-cortex, default-mode, and frontoparietal networks during bistable perception. Brain Connect. 2013 Jan;3(3):286–93.

115. Hayworth KJ, Lescroart MD, Biederman I. Neural encoding of relative position. J Exp Psychol Hum Percept Perform. 2011 Aug;37(4):1032–50.

116. Noudoost B, Chang MH, Steinmetz N a, Moore T. Top-down control of visual attention. Curr Opin Neurobiol. 2010 Apr;20(2):183–90.

117. Corbetta M, Kincade JM, Ollinger JM, McAvoy MP, Shulman GL. Voluntary orienting is dissociated from target detection in human posterior parietal cortex. Nat Neurosci. 2000 Mar;3(3):292–7.

118. Hopfinger J, Buonocore M, Mangun G. The neural mechanisms of top-down attentional control. Nat Neurosci. 2000;3:284–91.

119. Kastner S, Pinsk MA, De Weerd P, Desimone R, Ungerleider LG. Increased activity in human visual cortex during directed attention in the absence of visual stimulation. Neuron. 1999 Apr;22(4):751–61.

120. Pinto Y, van der Leij AR, Sligte IG, Lamme VAF, Scholte HS. Bottom-up and top-down attention are independent. J Vis. The Association for Research in Vision and Ophthalmology; 2013 Jan 1;13(3):16.

121. Heinen K, Feredoes E, Weiskopf N, Ruff CC, Driver J. Direct evidence for attention-dependent influences of the frontal eye-fields on feature-responsive visual cortex. Cereb Cortex. 2014 Nov 21;24(11):2815–21.

122. Armstrong KM, Chang MH, Moore T. Selection and maintenance of spatial information by frontal eye field neurons. J Neurosci. 2009 Dec 16;29(50):15621–9.

123. Bressler SL, Tang W, Sylvester CM, Shulman GL, Corbetta M. Top-down control of human visual cortex by frontal and parietal cortex in anticipatory visual spatial attention. J Neurosci. 2008 Oct 1;28(40):10056–61.

124. Corbetta M, Tansy AP, Stanley CM, Astafiev S V, Snyder AZ, Shulman GL. A functional MRI study of preparatory signals for spatial location and objects. Neuropsychologia. 2005 Jan;43(14):2041–56.

125. Vuilleumier P, Schwartz S, Verdon V, Maravita A, Hutton C, Husain M, et al. Abnormal attentional modulation of retinotopic cortex in parietal patients with spatial neglect. Curr Biol. 2008 Oct 14;18(19):1525–9.

126. Gillebert CR, Mantini D, Thijs V, Sunaert S, Dupont P, Vandenberghe R. Lesion evidence for the critical role of the intraparietal sulcus in spatial attention. Brain. 2011 Jun;134(Pt 6):1694–709.

127. Vandenberghe R, Gillebert CR. Parcellation of parietal cortex: convergence between lesion-symptom mapping and mapping of the intact functioning brain. Behav Brain Res. 2009 May 16;199(2):171–82.

128. Driver J, Blankenburg F, Bestmann S, Ruff CC. New approaches to the study of human brain networks underlying spatial attention and related processes. Exp brain Res. 2010 Oct;206(2):153–62.

129. Ruff CC, Bestmann S, Blankenburg F, Bjoertomt O, Josephs O, Weiskopf N, et

al. Distinct causal influences of parietal versus frontal areas on human visual cortex: evidence from concurrent TMS-fMRI. Cereb Cortex. 2008 Apr 1;18(4):817–27.

130. Thomas AG, Marrett S, Saad ZS, Ruff D a, Martin A, Bandettini P a. Functional but not structural changes associated with learning: an exploration of longitudinal voxel-based morphometry (VBM). Neuroimage. Elsevier B.V.; 2009 Oct 15;48(1):117–25.

131. Vossel S, Weidner R, Driver J, Friston KJ, Fink GR. Deconstructing the architecture of dorsal and ventral attention systems with dynamic causal modeling. J Neurosci. 2012 Aug 1;32(31):10637–48.

132. Ruff CC, Blankenburg F, Bjoertomt O, Bestmann S, Freeman E, Haynes J-D, et al. Concurrent TMS-fMRI and psychophysics reveal frontal influences on human retinotopic visual cortex. Curr Biol. 2006 Aug 8;16(15):1479–88.

133. Grent-'t-Jong T, Woldorff MG. Timing and sequence of brain activity in top-down control of visual-spatial attention. PLoS Biol. 2007 Jan 2;5(1):e12.

134. van Boxtel JJ a, Tsuchiya N, Koch C. Consciousness and attention: on sufficiency and necessity. Front Psychol. 2010 Jan;1(December):217.

135. Rees G, Kreiman G, Koch C. Neural correlates of consciousness in humans. Nat Rev Neurosci. 2002 Apr;3(4):261–70.

136. Crick F, Koch C. A framework for consciousness. 2003;6(2):119–26.

137. Esterman M, Yantis S. Perceptual expectation evokes category-selective cortical activity. Cereb Cortex. 2010 May;20(5):1245–53.

138. Summerfield C, Koechlin E. A neural representation of prior information during perceptual inference. Neuron. 2008 Jul 31;59(2):336–47.

139. Kanai R, Komura Y, Shipp S, Friston K, Komura Y, Shipp S. Cerebral hierarchies : predictive processing , precision and the pulvinar. 2015;

140. Feldman H, Friston KJ. Attention, uncertainty, and free-energy. Front Hum Neurosci. 2010 Jan;4(December):215.

141. Watanabe T, Masuda N, Megumi F, Kanai R, Rees G. Energy landscape and dynamics of brain activity during human bistable perception. Nat Commun. Nature Publishing Group; 2014 Jan;5:4765.

142. Panagiotaropoulos TI, Kapoor V, Logothetis NK. Subjective visual perception : from local processing to emergent phenomena of brain activity. 2014;

143. Meyer K. Primary sensory cortices, top-down projections and conscious experience. Prog Neurobiol. 2011 Sep 1;94(4):408–17.

144. Mechelli A, Price CJ, Friston KJ. Where Bottom-up Meets Top-down : Neuronal Interactions during Perception and Imagery. 2004;(November):1256–65.

145. Büchel C, Friston KJ. Modulation of Connectivity in Visual Pathways by Attention : Cortical Interactions Evaluated with Structural Equation Modelling and fMRI. 1997;768–78.

146. George N, Dolan RJ, Fink GR, Baylis GC, Russell C, Driver J. Contrast polarity and face recognition in the human fusiform gyrus. 1999;2(6).

147. Thyroid LN. How the brain learns to see objects and faces in an impoverished context. 1997;389(October):352–3.

148. Slotnick SD, Thompson WL, Kosslyn SM. Visual mental imagery induces retinotopically organized activation of early visual areas. Cereb Cortex. 2005 Oct;15(10):1570–83.

149. Slotnick William|Kosslyn,Stephen S. Visual memory and visual mental imagery recruit common control and sensory regions of the brain. Cogn Neurosci. 2012;3(1):14–20.

150. Hsieh P-J, Colas JT, Kanwisher NG. Pre-stimulus pattern of activity in the fusiform face area predicts face percepts during binocular rivalry. Neuropsychologia. Elsevier Ltd; 2012 Mar;50(4):522–9.

151. O'Craven KM, Kanwisher N. Mental imagery of faces and places activates corresponding stiimulus-specific brain regions. J Cogn Neurosci. 2000 Nov;12(6):1013–23.

152. Eger E, Henson RN, Driver J, Dolan RJ. Mechanisms of top-down facilitation in perception of visual objects studied by FMRI. Cereb Cortex. 2007 Sep;17(9):2123–33.

153. Dentico D, Leung B, Chang J, Guokas J, Boly M, Tononi G, et al. NeuroImage Reversal of cortical information fl ow during visual imagery as compared to

visual perception. 2014;100:237–43.

154. Mechelli A, Price CJ, Friston KJ, Ishai A. Where bottom-up meets top-down: neuronal interactions during perception and imagery. Cereb Cortex. 2004 Nov;14(11):1256–65.

155. Wang M, Arteaga D, He BJ. Brain mechanisms for simple perception and bistable perception. 2013;2013.

156. von Helmholtz H. Treatise on Physiological Optics. Treatise on Physiological Optics,Vol 3, ed J P Southall, English translation, NewYork: Dover. 1924.

157. Westheimer G. Was Helmholtz a Bayesian? Perception. 2008;37(5):642–50.

158. Gregory RL. Knowledge in perception and illusion. Philos Trans R Soc Lond B Biol Sci. 1997 Aug 29;352(1358):1121–7.

159. MacKay DM. The epistemological problem for automata. In: Shanon CE, McCarthy NJ, editors. Automata Studies. Princeton University Press; 1956. p. 235–51.

160. Dayan P, Hinton GE, Neal RM, Zemel RS. The Helmholtz machine. Neural Comput. 1995 Sep;7(5):889–904.

161. Ballard DH, Hinton GE, Sejnowski TJ. Parallel visual computation. Nature. 1983;306(5938):21–6.

162. Friston KJ, Stephan KE. Free-energy and the brain. Synthese. 2007 Dec 1;159(3):417–58.

163. Friston K. The free-energy principle: a rough guide to the brain? Trends Cogn Sci. 2009 Jul;13(7):293–301.

164. Kersten D, Mamassian P, Yuille A. Object Perception as Bayesian Inference. Annu Rev Psychol. Annual Reviews; 2004 Feb 12;55(1):271–304.

165. Knill DC, Pouget A. The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci. 2004 Dec;27(12):712–9.

166. Mumford D. On the computational architecture of the neocortex. Biol Cybern. 1992 Jan;66(3):241–51.

167. Murray SO, Schrater P, Kersten D. Perceptual grouping and the interactions between visual cortical areas. Neural Netw. 2004;17(5-6):695–705.

168. Rao RP, Ballard DH. Dynamic model of visual recognition predicts neural response properties in the visual cortex. Neural Comput. 1997 May 15;9(4):721–63.

169. Kawato M, Hayakawa H, Inui T. A forward-inverse optics model of reciprocal connections between visual cortical areas. Informa UK Ltd UK; 2009.

170. O'Craven KM, Downing PE, Kanwisher N. fMRI evidence for objects as the units of attentional selection. Nature. 1999 Oct 7;401(6753):584–7.

171. Tong F, Nakayama K, Vaughan JT, Kanwisher N. Binocular rivalry and visual awareness in human extrastriate cortex. Neuron. 1998 Oct;21(4):753–9.

172. Bloch F, Hansen WW, Packard M. Nuclear Induction. Phys Rev. 1946 Feb 1;69(3-4):127–127.

173. Purcell E, Torrey H, Pound R. Resonance Absorption by Nuclear Magnetic Moments in a Solid. Phys Rev. 1946 Jan;69(1-2):37–8.

174. Bloch F. Nuclear Induction. Phys Rev. 1946 Oct 1;70(7-8):460–74.

175. LAUTERBUR PC. Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance. Nature. 1973 Mar 16;242(5394):190–1.

176. Mansfield P. Multi-planar image formation using NMR spin echoes. J Phys C Solid State Phys. IOP Publishing; 1977 Feb 14;10(3):L55–8.

177. Damadian R, Goldsmith M, Minkoff L. NMR in cancer: XVI. FONAR image of the live human body. Physiol Chem Phys. 1977 Jan;9(1):97–100, 108.

178. Brown GG, Perthen JE, Liu TT, Buxton RB. A primer on functional magnetic resonance imaging. Neuropsychol Rev. 2007 Jun;17(2):107–25.

179. Plewes DB, Kucharczyk W. Physics of MRI: a primer. J Magn Reson Imaging. 2012 May;35(5):1038–54.

180. Liney G. MRI From A to Z A Definitive Guide for Medical Professionals. 2005.

181. Currie S, Hoggard N, Craven IJ, Hadjivassiliou M, Wilkinson ID. Understanding MRI: basic MR physics for physicians. Postgrad Med J. 2013 Apr;89(1050):209–23.

182. Logothetis NK. The neural basis of the blood-oxygen-level-dependent

functional magnetic resonance imaging signal. Philos Trans R Soc Lond B Biol Sci. 2002 Aug 29;357(1424):1003–37.

183. Buxton RB. The physics of functional magnetic resonance imaging (fMRI). Reports Prog Phys. 2013 Sep 1;76(9):096601.

184. Pauling L, Coryell CD. The Magnetic Properties and Structure of Hemoglobin, Oxyhemoglobin and Carbonmonoxyhemoglobin. Proc Natl Acad Sci U S A. 1936 Apr;22(4):210–6.

185. Boxerman JL, Hamberg LM, Rosen BR, Weisskoff RM. Mr contrast due to intravascular magnetic susceptibility perturbations. Magn Reson Med. 1995 Oct;34(4):555–66.

186. Ogawa S, Menon RS, Tank DW, Kim SG, Merkle H, Ellermann JM, et al. Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. Biophys J. 1993 Mar;64(3):803–12.

187. Boxerman JL, Bandettini PA, Kwong KK, Baker JR, Davis TL, Rosen BR, et al. The intravascular contribution to fmri signal change: monte carlo modeling and diffusion-weighted studiesin vivo. Magn Reson Med. 1995 Jul;34(1):4–10.

188. Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, Hyde JS. Time course EPI of human brain function during task activation. Magn Reson Med. 1992 Jun;25(2):390–7.

189. Buxton RB. Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques. Cambridge University Press; 2009. 457 p.

190. Heeger DJ, Ress D. What does fMRI tell us about neuronal activity? Nat Rev Neurosci. 2002 Feb;3(2):142–51.

191. Hall CN, Reynell C, Gesslein B, Hamilton NB, Mishra A, Sutherland B a, et al. Capillary pericytes regulate cerebral blood flow in health and disease. Nature. Nature Publishing Group; 2014 Apr 3;508(7494):55–60.

192. van Zijl PCM, Hua J, Lu H. The BOLD post-stimulus undershoot, one of the most debated issues in fMRI. Neuroimage. 2012 Aug 15;62(2):1092–102.

193. Mullinger KJ, Mayhew SD, Bagshaw AP, Bowtell R, Francis ST. Poststimulus undershoots in cerebral blood flow and BOLD fMRI responses are modulated

by poststimulus neuronal activity. Proc Natl Acad Sci U S A. 2013 Aug 13;110(33):13636–41.

194. Logothetis NK. What we can do and what we cannot do with fMRI. Nature. 2008 Jun 12;453(7197):869–78.

195. Buxton RB, Griffeth VEM, Simon AB, Moradi F, Shmuel A. Variability of the coupling of blood flow and oxygen metabolism responses in the brain: a problem for interpreting BOLD studies but potentially a new window on the underlying neural activity. Front Neurosci. 2014 Jan;8(June):139.

196. Fox P, Raichle M, Mintun M, Dence C. Nonoxidative glucose consumption during focal physiologic neural activity. Science (80-). 1988 Jul 22;241(4864):462–4.

197. Attwell D, Iadecola C. The neural basis of functional brain imaging signals. 2002;25(12):621–5.

198. Ekstrom A. How and when the fMRI BOLD signal relates to underlying neural activity: the danger in dissociation. Brain Res Rev. Elsevier B.V.; 2010 Mar;62(2):233–44.

199. Lindauer U, Dirnagl U, Füchtemeier M, Böttiger C, Offenhauser N, Leithner C, et al. Pathophysiological interference with neurovascular coupling - when imaging based on hemoglobin might go blind. Front Neuroenergetics. 2010 Jan;2.

200. Wolf T, Lindauer U, Villringer A, Dirnagl U. Excessive oxygen or glucose supply does not alter the blood flow response to somatosensory stimulation or spreading depression in rats. Brain Res. 1997 Jul;761(2):290–9.

201. Powers WJ, Hirsch IB, Cryer PE. Effect of stepped hypoglycemia on regional cerebral blood flow response to physiological brain activation. Am J Physiol. 1996 Mar;270(2 Pt 2):H554–9.

202. Yablonskiy DA, Ackerman JJH, Raichle ME. Coupling between changes in human brain temperature and oxidative metabolism during prolonged visual stimulation. Proc Natl Acad Sci. 2000 Jun 20;97(13):7603–8.

203. Devor A, Sakadzic S, Saisan PA, Yaseen MA, Roussakis E, Srinivasan VJ, et al. “Overshoot” of O_2 is required to maintain baseline tissue oxygenation at

locations distal to blood vessels. J Neurosci. 2011 Sep 21;31(38):13676–81.

204. McCaslin AFH, Chen BR, Radosevich AJ, Cauli B, Hillman EMC. In vivo 3D morphology of astrocyte-vasculature interactions in the somatosensory cortex: implications for neurovascular coupling. J Cereb Blood Flow Metab. 2011 Mar;31(3):795–806.

205. Sun W, McConnell E, Pare J-F, Xu Q, Chen M, Peng W, et al. Glutamate-dependent neuroglial calcium signaling differs between young and adult brain. Science. 2013 Jan 11;339(6116):197–200.

206. Nizar K, Uhlirova H, Tian P, Saisan PA, Cheng Q, Reznichenko L, et al. In vivo stimulus-induced vasodilation occurs without IP3 receptor activation and may precede astrocytic calcium increase. J Neurosci. 2013 May 8;33(19):8411–22.

207. Takata N, Nagai T, Ozawa K, Oe Y, Mikoshiba K, Hirase H. Cerebral blood flow modulation by Basal forebrain or whisker stimulation can occur independently of large cytosolic Ca2+ signaling in astrocytes. PLoS One. 2013 Jan;8(6):e66525.

208. Dore-Duffy P, Cleary K. Morphology and properties of pericytes. Methods Mol Biol. 2011 Jan;686:49–68.

209. Peppiatt CM, Howarth C, Mobbs P, Attwell D. Bidirectional control of CNS capillary diameter by pericytes. Nature. 2006 Oct 12;443(7112):700–4.

210. Puro DG. Physiology and pathobiology of the pericyte-containing retinal microvasculature: new developments. Microcirculation. 2007 Jan;14(1):1–10.

211. Bagher P, Segal SS. Regulation of blood flow in the microcirculation: role of conducted vasodilation. Acta Physiol (Oxf). 2011 Jul;202(3):271–84.

212. Figueroa XF, Duling BR. Gap junctions in the control of vascular function. Antioxid Redox Signal. 2009 Mar;11(2):251–66.

213. Chen BR, Bouchard MB, McCaslin AFH, Burgess SA, Hillman EMC. High-speed vascular dynamics of the hemodynamic response. Neuroimage. 2011 Jan 15;54(2):1021–30.

214. Chen BR, Kozberg MG, Bouchard MB, Shaik MA, Hillman EMC. A critical role for the vascular endothelium in functional neurovascular coupling in the brain. J Am Heart Assoc. 2014 Jun 12;3(3):e000787.

215. Tian P, Teng IC, May LD, Kurz R, Lu K, Scadeng M, et al. Cortical depth-specific microvascular dilation underlies laminar differences in blood oxygenation level-dependent functional MRI signal. Proc Natl Acad Sci U S A. 2010 Aug 24;107(34):15246–51.

216. Arnerić SP, Honig MA, Milner TA, Greco S, Iadecola C, Reis DJ. Neuronal and endothelial sites of acetylcholine synthesis and release associated with microvessels in rat cerebral cortex: ultrastructural and neurochemical studies. Brain Res. 1988 Jun 28;454(1-2):11–30.

217. Hotta H, Uchida S, Kagitani F, Maruyama N. Control of cerebral cortical blood flow by stimulation of basal forebrain cholinergic areas in mice. J Physiol Sci. 2011 May;61(3):201–9.

218. Kocharyan A, Fernandes P, Tong X-K, Vaucher E, Hamel E. Specific subtypes of cortical GABA interneurons contribute to the neurovascular coupling response to basal forebrain stimulation. J Cereb Blood Flow Metab. 2008 Mar;28(2):221–31.

219. Piché M, Uchida S, Hara S, Aikawa Y, Hotta H. Modulation of somatosensory-evoked cortical blood flow changes by GABAergic inhibition of the nucleus basalis of Meynert in urethane-anaesthetized rats. J Physiol. 2010 Jun 15;588(Pt 12):2163–71.

220. Sato A, Sato Y. Cholinergic neural regulation of regional cerebral blood flow. Alzheimer Dis Assoc Disord. 1995 Jan;9(1):28–38.

221. Cauli B, Hamel E. Revisiting the role of neurons in neurovascular coupling. Front Neuroenergetics. 2010 Jan;2:9.

222. Cauli B, Tong X-K, Rancillac A, Serluca N, Lambolez B, Rossier J, et al. Cortical GABA interneurons in neurovascular coupling: relays for subcortical vasoactive pathways. J Neurosci. 2004 Oct 13;24(41):8940–9.

223. Bekar LK, Wei HS, Nedergaard M. The locus coeruleus-norepinephrine network optimizes coupling of cerebral blood volume with oxygen demand. J Cereb Blood Flow Metab. 2012 Dec;32(12):2135–45.

224. Kozberg MG, Chen BR, DeLeo SE, Bouchard MB, Hillman EMC. Resolving the transition from negative to positive blood oxygen level-dependent

responses in the developing brain. Proc Natl Acad Sci U S A. 2013 Mar 12;110(11):4380–5.

225. Hillman EMC. Coupling Mechanism and Significance of the BOLD Signal: A Status Report. Annu Rev Neurosci. 2014 Jul 8;37:161–81.

226. Logothetis NK, Wandell B a. Interpreting the BOLD signal. Annu Rev Physiol. 2004 Jan;66:735–69.

227. Nair DG. About being BOLD. Brain Res Brain Res Rev. 2005 Dec 15;50(2):229–43.

228. Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann a. Neurophysiological investigation of the basis of the fMRI signal. Nature. 2001 Jul 12;412(6843):150–7.

229. Herman P, Sanganahalli BG, Blumenfeld H, Rothman DL, Hyder F. Quantitative basis for neuroimaging of cortical laminae with calibrated functional MRI. Proc Natl Acad Sci U S A. 2013 Sep 10;110(37):15115–20.

230. Kim DS, Duong TQ, Kim SG. High-resolution mapping of iso-orientation columns by fMRI. Nat Neurosci. 2000 Mar;3(2):164–9.

231. Siero JCW, Hermes D, Hoogduin H, Luijten PR, Ramsey NF, Petridou N. BOLD matches neuronal activity at the mm scale: a combined 7T fMRI and ECoG study in human sensorimotor cortex. Neuroimage. Elsevier Inc.; 2014 Nov 1;101:177–84.

232. Weiskopf N, Veit R, Erb M, Mathiak K, Grodd W, Goebel R, et al. Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. Neuroimage. 2003 Jul;19(3):577–86.

233. Zatorre RJ, Fields RD, Johansen-Berg H. Plasticity in gray and white: neuroimaging changes in brain structure during learning. Nat Neurosci. Nature Publishing Group; 2012 Apr;15(4):528–36.

234. Lindquist M a. The Statistical Analysis of fMRI Data. Stat Sci. 2008 Nov;23(4):439–64.

235. Ashby FG. Statistical Analysis of FMRI Data. MIT Press; 2011. 332 p.

236. Friston KJ, Frith CD, Turner R, Frackowiak RSJ. Characterizing Evoked

Hemodynamics with fMRI. Neuroimage. 1995 Jun;2(2):157–65.

237. Heuttel SA, Song AW, G M. Functional Magnetic Resonance Imaging. 2004.

238. Friston KJ, Ashburner J, Frith CD, Poline J-B, Heather JD, Frackowiak RSJ. Spatial registration and normalization of images. Hum Brain Mapp. 1995 Oct 13;3(3):165–89.

239. Friston KJ, Holmes AP, Worsley KJ, Poline J-P, Frith CD, Frackowiak RSJ. Statistical parametric maps in functional imaging: A general linear approach. Hum Brain Mapp. 1994 Oct 13;2(4):189–210.

240. Hutton C, Bork A, Josephs O, Deichmann R, Ashburner J, Turner R. Image distortion correction in fMRI: A quantitative evaluation. Neuroimage. 2002 May;16(1):217–40.

241. Evans AC, Collins DL, Mills SR, Brown ED, Kelly RL, Peters TM. 3D statistical neuroanatomical models from 305 MRI volumes. 1993 IEEE Conference Record Nuclear Science Symposium and Medical Imaging Conference. IEEE; 1993. p. 1813–7.

242. Talairach J, Tournoux P. Co-planar stereotaxic atlas of the human brain. 3-Dimensional proportional system: an approach to cerebral imaging. 1988;

243. Li X, Yao L, Ye Q, Zhao X. Online spatial normalization for real-time FMRI. PLoS One. 2014 Jan;9(7):e103302.

244. Ashburner J, Friston KJ. Nonlinear spatial normalization using basis functions. Hum Brain Mapp. 1999;7(4):254–66.

245. Murphy KP. Machine learning, a probabilistic perspective. 2012.

246. Friston KJ, Fletcher P, Josephs O, Holmes A, Rugg MD, Turner R. Event-related fMRI: characterizing differential responses. Neuroimage. 1998 Jan;7(1):30–40.

247. Friston KJ, Josephs O, Rees G, Turner R. Nonlinear event-related responses in fMRI. Magn Reson Med. 1998 Jan;39(1):41–52.

248. Friston KJ, Penny W, Phillips C, Kiebel S, Hinton G, Ashburner J. Classical and Bayesian inference in neuroimaging: theory. Neuroimage. 2002 Jun;16(2):465–83.

249. Buračas GT, Boynton GM. Efficient Design of Event-Related fMRI Experiments Using M-Sequences. Neuroimage. 2002 Jul;16(3):801–13.

250. Jezzard P, Matthews PM, Smith SM. Functional Magnetic Resonance Imaging: Peter Jezzard - Oxford University Press. 2003.

251. Worsley KJ, Taylor JE. Detecting fMRI activation allowing for unknown latency of the hemodynamic response. Neuroimage. United States; 2006 Jan 15;29(2):649–54.

252. Friston KJ, Holmes AP, Worsley KJ. How many subjects constitute a study? Neuroimage. 1999 Jul;10(1):1–5.

253. Friston K. Ten ironic rules for non-statistical reviewers. Neuroimage. 2012 Jul 16;61(4):1300–10.

254. Lindquist MA, Caffo B, Crainiceanu C. Ironing out the statistical wrinkles in “ten ironic rules”. Neuroimage. 2013 Nov 1;81:499–502.

255. Holmes A, Friston KJ. Generalisability, Random Effects and Population Inference. Neuroimage. 1998;7:S754.

256. Haxby J V, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science. 2001 Sep 28;293(5539):2425–30.

257. Cox DD, Savoy RL. Functional magnetic resonance imaging (fMRI) “brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex. Neuroimage. 2003 Jun;19(2):261–70.

258. Hanson S, Matsuka T, Haxby J. Combinatorial codes in ventral temporal lobe for object recognition: Haxby(2001) revisited- is there a face area. Neuroimage. 2004 Jan 1;23:163–6.

259. Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain. Nat Neurosci. 2005 May;8(5):679–85.

260. Haynes J-D, Deichmann R, Rees G. Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. Nature. 2005 Nov 24;438(7067):496–9.

261. Kriegeskorte N, Goebel R, Bandettini P. Information-based functional brain mapping. Proc Natl Acad Sci U S A. 2006 Mar 7;103(10):3863–8.

262. Kriegeskorte N, Bandettini P. Combining the tools: activation-and information-based fMRI analysis. Neuroimage. 2007;

263. Quian Quiroga R, Panzeri S. Extracting information from neuronal populations: information theory and decoding approaches. Nat Rev Neurosci. Nature Publishing Group; 2009 Mar;10(3):173–85.

264. Norman KA, Polyn SM, Detre GJ, Haxby J V. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends Cogn Sci. 2006 Sep;10(9):424–30.

265. Kriegeskorte N. Pattern-information analysis: from stimulus decoding to computational-model testing. Neuroimage. Elsevier Inc.; 2011 May 15;56(2):411–21.

266. Pereira F, Mitchell T, Botvinick M. Machine learning classifiers and fMRI: a tutorial overview. Neuroimage. Elsevier Inc.; 2009 Mar;45(1 Suppl):S199–209.

267. Samuel Schwarzkopf D, Rees G. Pattern classification using functional magnetic resonance imaging. Wiley Interdiscip Rev Cogn Sci. 2011 Sep;2(5):568–79.

268. Kohavi R. A study of cross-validation and bootstrap for accuracy estimation and model selection. Morgan Kaufmann Publishers Inc.; 1995 Aug 20;1137–43.

269. Misaki M, Kim Y, Bandettini P a, Kriegeskorte N. Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. Neuroimage. Elsevier B.V.; 2010 Oct 15;53(1):103–18.

270. Strother SC, Anderson J, Hansen LK, Kjems U, Kustra R, Sidtis J, et al. The quantitative evaluation of functional neuroimaging experiments: the NPAIRS data analysis framework. Neuroimage. 2002 Apr;15(4):747–71.

271. Duda R., Hart PE, Stork D. Pattern Classification. Journal of Classification. New York: John Wiley & Sons; 2001.

272. Pereira F, Mitchell T, Botvinick M. Machine learning classifiers and fMRI : A tutorial overview. Neuroimage. Elsevier Inc.; 2009;45(1):S199–209.

273. LaConte SM. Decoding fMRI brain states in real-time. Neuroimage. Elsevier B.V.; 2011 May 15;56(2):440–54.

274. Bahramisharif A, Gerven M Van, Heskes T, Jensen O. Covert attention allows

for continuous control of brain – computer interfaces. 2010;31:1501–8.

275. Fitzgibbon T, Taylor SF. Retinotopy of the human retinal nerve fibre layer and optic nerve head. J Comp Neurol. 1996 Nov 11;375(2):238–51.

276. Horton J, Greenwood M, Hubel D. Non-retinotopic arrangement of fibres in cat optic nerve. Nature. 1979;

277. HENSCHEN SE. ON THE VISUAL PATH AND CENTRE. Brain. 1893 Jan 1;16(1-2):170–80.

278. Inouye T. Die Sehstörungen bei Schussverletzungen der kortikalen Sehsphäre: nach Beobachtungen an Verwundeten der letzten japanischen Kriege. 1909;

279. HOLMES G, LISTER WT. DISTURBANCES OF VISION FROM CEREBRAL LESIONS, WITH SPECIAL REFERENCE TO THE CORTICAL REPRESENTATION OF THE MACULA. Brain. 1916 Jun 1;39(1-2):34–73.

280. Holmes G. DISTURBANCES OF VISUAL ORIENTATION. Br J Ophthalmol. 1918 Sep;2(9):449–68.

281. Horton JC, Hoyt WF. The Representation of the Visual Field in Human Striate Cortex- a revision of the classic holmes map. Arch Ophthalmol. American Medical Association; 1991 Jun 1;109(6):816.

282. Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, et al. fMRI of human visual cortex. Nature. ENGLAND; 1994 Jun 16;369(6481):525.

283. Engel S a, Glover GH, Wandell BA. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cereb Cortex. 1997;7(2):181–92.

284. Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, et al. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science. UNITED STATES; 1995 May 12;268(5212):889–93.

285. Dale AM, Fischl B, Sereno MI. Cortical surface-based analysis. I. Segmentation and surface reconstruction. Neuroimage. 1999 Feb;9(2):179–94.

286. Ségonne F, Dale AM, Busa E, Glessner M, Salat D, Hahn HK, et al. A hybrid approach to the skull stripping problem in MRI. Neuroimage. 2004 Jul;22(3):1060–75.

287. Fischl B, Salat DH, Busa E, Albert M, Dieterich M, Haselgrove C, et al. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron. 2002 Jan 31;33(3):341–55.

288. Fischl B, Salat DH, van der Kouwe AJW, Makris N, Ségonne F, Quinn BT, et al. Sequence-independent segmentation of magnetic resonance images. Neuroimage. 2004 Jan;23 Suppl 1:S69–84.

289. Sled JG, Zijdenbos AP, Evans AC. A nonparametric method for automatic correction of intensity nonuniformity in MRI data. IEEE Trans Med Imaging. 1998 Feb;17(1):87–97.

290. Fischl B, Liu A, Dale AM. Automated manifold surgery: constructing geometrically accurate and topologically correct models of the human cerebral cortex. IEEE Trans Med Imaging. IEEE; 2001 Jan 1;20(1):70–80.

291. Ségonne F, Pacheco J, Fischl B. Geometrically accurate topology-correction of cortical surfaces using nonseparating loops. IEEE Trans Med Imaging. IEEE; 2007 Apr 1;26(4):518–29.

292. Fischl B, Sereno MI, Dale AM. Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. Neuroimage. 1999 Feb;9(2):195–207.

293. Fischl B, Dale AM. Measuring the thickness of the human cerebral cortex from magnetic resonance images. Proc Natl Acad Sci U S A. 2000 Sep 26;97(20):11050–5.

294. Ashburner J. A fast diffeomorphic image registration algorithm. Neuroimage. 2007 Oct 15;38(1):95–113.

295. Li W, He H, Lu J, Lv B, Li M, Jin Z. <title>Detection of whole-brain abnormalities in temporal lobe epilepsy using tensor-based morphometry with DARTEL</title>. Liu J, Doi K, Fenster A, Chan SC, editors. 2009 Oct 30;7497:749723–749723 – 6.

296. Ceccarelli A, Rocca MA, Pagani E, Falini A, Comi G, Filippi M. Cognitive

learning is associated with gray matter changes in healthy human individuals: a tensor-based morphometry study. Neuroimage. Elsevier Inc.; 2009 Nov 15;48(3):585–9.

297. Welch K a, Moorhead TW, McIntosh a M, Owens DGC, Johnstone EC, Lawrie SM. Tensor-based morphometry of cannabis use on brain structure in individuals at elevated genetic risk of schizophrenia. Psychol Med. 2013 Oct;43(10):2087–96.

298. Wang Y, Yuan L, Shi J, Greve A, Ye J, Toga AW, et al. Applying tensor-based morphometry to parametric surfaces can improve MRI-based disease diagnosis. Neuroimage. Elsevier Inc.; 2013 Jul 1;74:209–30.

299. Farbota KDM, Sodhi A, Bendlin BB, McLaren DG, Xu G, Rowley H a, et al. Longitudinal volumetric changes following traumatic brain injury: a tensor-based morphometry study. J Int Neuropsychol Soc. 2012 Nov;18(6):1006–18.

300. Goebel R, Esposito F, Formisano E. Analysis of functional image analysis contest (FIAC) data with brainvoyager QX: From single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. Hum Brain Mapp. 2006 May;27(5):392–401.

301. Weiskopf N, Mathiak K, Bock SW, Scharnowski F, Veit R, Grodd W, et al. Principles of a brain-computer interface (BCI) based on real-time functional magnetic resonance imaging (fMRI). IEEE Trans Biomed Eng. 2004 Jun;51(6):966–70.

302. Koush Y, Zvyagintsev M, Dyck M, Mathiak K a, Mathiak K. Signal quality and Bayesian signal processing in neurofeedback based on real-time fMRI. Neuroimage. Elsevier Inc.; 2012 Jan 2;59(1):478–89.

303. Johnson K a, Hartwell K, LeMatty T, Borckardt J, Morgan PS, Govindarajan K, et al. Intermittent “real-time” fMRI feedback is superior to continuous presentation for a motor imagery task: a pilot study. J Neuroimaging. 2012 Jan;22(1):58–66.

304. Sarkheil P, Zilverstand A, Kilian-Huetten, N Schneider F, Mathiak K, Goebel R. Application of real-timefMRI-neurofeedback in treatmentof emotional disorders. Annual Meeting of the OHBM, Quebec City, Canada. Quebec City, Canada; 2011.

305. Megumi F, Yamashita A, Kawato M, Imamizu H. Functional MRI neurofeedback training on connectivity between two regions induces long-lasting changes in intrinsic functional network. Front Hum Neurosci. 2015;9(March).

306. Bagarinao E, Matsuo K, Nakai T, Sato S. Estimation of general linear model coefficients for real-time application. Neuroimage. 2003 Jun;19(2):422–9.

307. Bagarinao E, Matsuo K, Tanaka Y, Sarmenta LFG, Nakai T. Enabling On-demand Real-time Functional MRI Analysis Using Grid Technology. Methods Inf Med. 2005 Jan 12;44(5):665–73.

308. Zilverstand A, Sorger B, Zimmermann J, Kaas A, Goebel R. Windowed correlation: a suitable tool for providing dynamic fMRI-based functional connectivity neurofeedback on task difficulty. PLoS One. 2014 Jan;9(1):e85929.

309. Posse S, Fitzgerald D, Gao K, Habel U, Rosenberg D, Moore GJ, et al. Real-time fMRI of temporolimbic regions detects amygdala activation during single-trial self-induced sadness. Neuroimage. 2003 Mar;18(3):760–8.

310. Hinds O, Ghosh S, Thompson TW, Yoo JJ, Whitfield-Gabrieli S, Triantafyllou C, et al. Computing moment-to-moment BOLD activation for real-time neurofeedback. Neuroimage. Elsevier Inc.; 2011 Jan 1;54(1):361–8.

311. Andersson P, Viergever M a., Pluim JPW, Ramsey NF, Siero JCW. fMRI based BCI control using spatial visual attention at 7T. 2009 4th Int IEEE/EMBS Conf Neural Eng. Ieee; 2009 Apr;444–6.

312. Weiskopf N. Neurofeedback training in Huntington's disease: enhancing neural plasticity using real-time fMRI neurofeedback training. Real-time Functional Imaging and Neurofeedback. 2015.

313. Gao K, Posse S. TurboFIRE: Real-time fMRI with automated spatial normalization and Talairach Daemon database. New York, USA. 9th Annual Meeting of the OHBM,. New York; 2003.

314. Zotev V, Phillips R, Yuan H, Misaki M, Bodurka J. Self-regulation of human brain activity using simultaneous real-time fMRI and EEG neurofeedback. Neuroimage. 2014 Jan 15;85 Pt 3:985–95.

315. Caria A, Veit R, Sitaram R, Lotze M, Weiskopf N, Grodd W, et al. Regulation of anterior insular cortex activity using real-time fMRI. Neuroimage. 2007 Apr 15;35(3):1238–46.

316. Caria A, Sitaram R, Veit R, Begliomini C, Birbaumer N. Volitional control of anterior insula activity modulates the response to aversive stimuli. A real-time functional magnetic resonance imaging study. Biol Psychiatry. Elsevier Inc.; 2010 Sep 1;68(5):425–32.

317. Berman BD, Horovitz SG, Venkataraman G, Hallett M. Self-modulation of primary motor cortex activity with motor and motor imagery tasks using real-time fMRI-based neurofeedback. Neuroimage. Elsevier Inc.; 2012 Jan 16;59(2):917–25.

318. Chiew M, LaConte SM, Graham SJ. Investigation of fMRI neurofeedback of differential primary motor cortex activity using kinesthetic motor imagery. Neuroimage. 2012;61(1):21–31.

319. Yoo S-S, O'Leary HM, Lee J-H, Chen N-K, Panych LP, Jolesz F a. Reproducibility of trial-based functional MRI on motor imagery. Int J Neurosci. 2007 Mar;117(2):215–27.

320. Yoo SS, O'Leary HM, Fairneny T, Chen NK, Panych LP, Park H, et al. Increasing cortical activity in auditory areas through neurofeedback functional magnetic resonance imaging. Neuroreport. 2006;17(12):1273–8.

321. Saxe R, Brett M, Kanwisher N. Divide and conquer: a defense of functional localizers. Neuroimage. 2006 May 1;30(4):1088–96; discussion 1097–9.

322. Sulzer J, Haller S, Scharnowski F, Weiskopf N, Birbaumer N, Blefari ML, et al. Real-time fMRI neurofeedback: progress and challenges. Neuroimage. 2013 Aug 1;76:386–99.

323. deBettencourt MT, Cohen JD, Lee RF, Norman K a, Turk-Browne NB. Closed-loop training of attention with real-time brain imaging. Nat Neurosci. Nature Publishing Group; 2015 Feb 9;(October 2014):1–9.

324. Stephan KM, Fink GR, Passingham RE, Silbersweig D, Ceballos-Baumann AO, Frith CD, et al. Functional anatomy of the mental representation of upper extremity movements in healthy subjects. J Neurophysiol. 1995

Jan;73(1):373–86.

325. Porro CA, Francescato MP, Cettolo V, Diamond ME, Baraldi P, Zuiani C, et al. Primary motor and sensory cortex activation during motor performance and motor imagery: a functional magnetic resonance imaging study. J Neurosci. 1996 Dec 1;16(23):7688–98.

326. Veit R, Singh V, Sitaram R, Caria A, Rauss K, Birbaumer N. Using real-time fMRI to learn voluntary regulation of the anterior insula in the presence of threat-related stimuli. Soc Cogn Affect Neurosci. 2012 Aug 1;7(6):623–34.

327. Bray S, Chang C, Hoeft F. Applications of multivariate pattern classification analyses in developmental neuroimaging of healthy and clinical populations. Front Hum Neurosci. 2009;3(32).

328. Yoo JJ, Hinds O, Ofen N, Thompson TW, Whitfield-Gabrieli S, Triantafyllou C, et al. When the brain is prepared to learn: enhancing human learning using real-time fMRI. Neuroimage. Elsevier Inc.; 2012 Jan 2;59(1):846–52.

329. Fox SS, Rudell a P. Operant controlled neural event: formal and systematic approach to electrical coding of behavior in brain. Science. 1968 Dec 13;162(3859):1299–302.

330. Sitaram R, Veit R, Stevens B, Caria A, Gerloff C, Birbaumer N, et al. Acquired control of ventral premotor cortex activity by feedback training: an exploratory real-time FMRI and TMS study. Neurorehabil Neural Repair. 2012;26(3):256–65.

331. Gulati T, Ramanathan DS, Wong CC, Ganguly K. Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. Nat Neurosci. Nature Publishing Group; 2014 Aug;17(8):1107–13.

332. Auer T, Schweizer R, Frahm J. Training Efficiency and Transfer Success in an Extended Real-Time Functional MRI Neurofeedback Training of the Somatomotor Cortex of Healthy Subjects. Front Hum Neurosci. 2015;9(October):1–14.

333. Ebbinghaus H. Memory: A Contribution to Experimental Psychology. University Microfilms; 1913. 123 p.

334. Hammer EM, Halder S, Blankertz B, Sannelli C, Dickhaus T, Kleih S, et al.

Psychological predictors of SMR-BCI performance. Biol Psychol. Elsevier B.V.; 2012 Jan;89(1):80–6.

335. Sulzer J, Haller S, Scharnowski F, Weiskopf N, Birbaumer N, Blefari ML, et al. Real-time fMRI neurofeedback: Progress and challenges. Neuroimage. Elsevier Inc.; 2013 Aug 1;76:386–99.

336. Egner T, Gruzelier JH. Learned self-regulation of EEG frequency components affects attention and event-related brain potentials in humans. Neuroreport. 2001;12(18):4155–9.

337. Reichert C, Fendrich R, Bernarding J, Tempelmann C, Hinrichs H, Rieger JW. Online tracking of the contents of conscious perception using real-time fMRI. Front Neurosci. 2014 Jan;8(May):116.

338. Hinterberger T, Schmidt S, Neumann N, Mellinger J, Blankertz B, Curio G, et al. Brain-computer communication and slow cortical potentials. IEEE Trans Biomed Eng. 2004;51(6):1011–8.

339. Ruiz S. Acquired self-control of insula cortex modulates emotion recognition and brain network connectivity in schizophrenia. Hum Brain Mapp. 2013;34(1):200–12.

340. Sorger B, Peters J, Boomen C Van Den, Zilverstand A, Reithler J, Goebel R. Real-time decoding of the locus of visuospatial attention using multi-voxel pattern classification. 2009;(2007):2007.

341. Weiskopf N, Scharnowski F, Veit R, Goebel R, Birbaumer N, Mathiak K. Self-regulation of local brain activity using real-time functional magnetic resonance imaging (fMRI). J Physiol Paris. 2005;98(4-6):357–73.

342. Sitaram R, Caria A, Veit R, Gaber T, Rota G, Kuebler A, et al. fMRI Brain-Computer Interface : A Tool for Neuroscientific Research and Treatment. 2007;2007.

343. Mueller C, Luehrs M, Baecke S, Adolf D, Luetzkendorf R, Luchtmann M, et al. Building virtual reality fMRI paradigms: a framework for presenting immersive virtual environments. J Neurosci Methods. Elsevier B.V.; 2012 Aug 15;209(2):290–8.

344. Cohen O, Koppel M, Malach R, Friedman D. Controlling an avatar by thought

using real-time fMRI. J Neural Eng. 2014 Jun;11(3):035006.

345. Hinterberger T, Weiskopf N, Veit R, Wilhelm B, Betta E, Birbaumer N. An EEG-driven brain-computer interface combined with functional magnetic resonance imaging (fMRI). IEEE Trans Biomed Eng. 2004;51(6):971–4.

346. Travis TA, Kondo CY, Knott JR. Parameters of eyes-closed alpha enhancement. Psychophysiology. UNITED STATES; 1974 Nov;11(6):674–81.

347. Berman BD, Horovitz SG, Hallett M. Modulation of functionally localized right insular cortex activity using real-time fMRI-based neurofeedback. Front Hum Neurosci. 2013 Jan;7:638.

348. Greer SM, Trujillo AJ, Glover GH, Knutson B. Control of nucleus accumbens activity with neurofeedback. Neuroimage. 2014 Aug 1;96:237–44.

349. Klingberg T. Training and plasticity of working memory. Trends Cogn Sci. Elsevier; 2010 Jul 7;14(7):317–24.

350. Ptak R, Schnider A. The attention network of the human brain: relating structural damage associated with spatial neglect to functional imaging correlates of spatial attention. Neuropsychologia. Elsevier Ltd; 2011 Sep;49(11):3063–70.

351. Adair JC, Barrett AM. Spatial neglect: clinical and neuroscience review: a wealth of information on the poverty of spatial attention. Ann N Y Acad Sci. 2008 Oct;1142:21–43.

352. Husain M, Nachev P. Space and the parietal cortex. Trends Cogn Sci. 2007 Jan;11(1):30–6.

353. Bartolomeo P, Broca CP. THE RELATIONSHIP BETWEEN VISUAL PERCEPTION AND VISUAL MENTAL IMAGERY : A REAPPRAISAL OF THE NEUROPSYCHOLOGICAL EVIDENCE. 2002;357–78.

354. Mort DJ, Malhotra P, Mannan SK, Rorden C, Pambakian A, Kennard C, et al. The anatomy of visual neglect. Brain. 2003 Sep;126(Pt 9):1986–97.

355. Golay L, Hauert C-A, Greber C, Schnider A, Ptak R. Dynamic modulation of visual detection by auditory cues in spatial neglect. Neuropsychologia. 2005 Jan;43(9):1258–65.

356. Golay L, Schnider A, Ptak R. Cortical and subcortical anatomy of chronic

spatial neglect following vascular damage. Behav Brain Funct. 2008 Jan;4(1):43.

357. Posner MI. Orienting of attention. Q J Exp Psychol. 1980 Feb;32(1):3–25.

358. William J. The Principles of Psychology [Internet]. 1890 [cited 2015 May 1]. Available from: https://ebooks.adelaide.edu.au/j/james/william/principles/

359. Goebel R, Zilverstand A, Sorger B. Real-time fMRI-based brain–computer interfacing for neurofeedback therapy and compensation of lost motor functions. Imaging Med. 2010 Aug;2(4):407–15.

360. Shipp S. The brain circuitry of attention. Trends Cogn Sci. 2004 May;8(5):223–30.

361. Bartolomeo P, Chokron S. Orienting of attention in left unilateral neglect. Neurosci Biobehav Rev. 2002 Mar;26(2):217–34.

362. Vossel S, Geng JJ, Fink GR. Dorsal and ventral attention systems: distinct neural circuits but collaborative roles. Neuroscientist. 2014 Apr;20(2):150–9.

363. DeYoe EA, Carman GJ, Bandettini P, Glickman S, Wieser J, Cox R, et al. Mapping striate and extrastriate visual areas in human cerebral cortex. Proc Natl Acad Sci U S A. 1996 Mar 19;93(6):2382–6.

364. Press WA, Brewer AA, Dougherty RF, Wade AR, Wandell BA. Visual areas and spatial summation in human visual cortex. Vision Res. 2001 Jan;41(10-11):1321–32.

365. Gattass R, Nascimento-Silva S, Soares JGM, Lima B, Jansen AK, Diogo ACM, et al. Cortical visual areas in monkeys: location, topography, connections, columns, plasticity and cortical dynamics. Philos Trans R Soc Lond B Biol Sci. 2005 Apr 29;360(1456):709–31.

366. Serences JT, Yantis S. Spatially selective representations of voluntary and stimulus-driven attentional priority in human occipital, parietal, and frontal cortex. Cereb Cortex. 2007 Feb 1;17(2):284–93.

367. Sclar G, Maunsell JH, Lennie P. Coding of image contrast in central visual pathways of the macaque monkey. Vision Res. 1990 Jan;30(1):1–10.

368. Sereno MI, Pitzalis S, Martinez A. Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. Science. 2001 Nov

9;294(5545):1350–4.

369. Huk AC, Dougherty RF, Heeger DJ. Retinotopy and functional subdivision of human areas MT and MST. J Neurosci. 2002;22(16):7195–205.

370. Hasson U, Harel M, Levy I, Malach R. Large-Scale Mirror-Symmetry Organization of Human Occipito-Temporal Object Areas. Neuron. 2003 Mar;37(6):1027–41.

371. Brewer AA, Liu J, Wade AR, Wandell BA. Visual field maps and stimulus selectivity in human ventral occipital cortex. Nat Neurosci. Nature Publishing Group; 2005 Aug;8(8):1102–9.

372. Schluppeck D, Glimcher P, Heeger DJ. Topographic organization for delayed saccades in human posterior parietal cortex. J Neurophysiol. 2005 Aug;94(2):1372–84.

373. Silver M a, Kastner S. Topographic maps in human frontal and parietal cortex. Trends Cogn Sci. 2009 Nov;13(11):488–95.

374. Hagler DJ, Sereno MI. Spatial maps in frontal and prefrontal cortex. Neuroimage. 2006 Jan 15;29(2):567–77.

375. Sereno MI, Huang R-S. A human parietal face area contains aligned head-centered visual and tactile maps. Nat Neurosci. Nature Publishing Group; 2006 Oct;9(10):1337–43.

376. Hagler DJ, Riecke L, Sereno MI. Parietal and superior frontal visuospatial maps activated by pointing and saccades. Neuroimage. 2007 May 1;35(4):1562–77.

377. Kastner S, DeSimone K, Konen CS, Szczepanski SM, Weiner KS, Schneider KA. Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. J Neurophysiol. 2007 May;97(5):3494–507.

378. Swisher JD, Halko MA, Merabet LB, McMains SA, Somers DC. Visual topography of human intraparietal sulcus. J Neurosci. 2007 May 16;27(20):5326–37.

379. Arcaro MJ, McMains S a, Singer BD, Kastner S. Retinotopic organization of human ventral visual cortex. J Neurosci. 2009 Aug 26;29(34):10638–52.

380. Malach R, Reppas JB, Benson RR, Kwong KK, Jiang H, Kennedy WA, et al. Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. Proc Natl Acad Sci U S A. 1995 Aug 29;92(18):8135–9.

381. Grill-Spector K, Kourtzi Z, Kanwisher N. The lateral occipital complex and its role in object recognition. Vision Res. 2001 Jan;41(10-11):1409–22.

382. Liu J, Harris A, Kanwisher N. Perception of face parts and face configurations: an FMRI study. J Cogn Neurosci. 2010 Jan;22(1):203–11.

383. Kaido T, Hoshida T, Taoka T, Sakaki T. Retinotopy with coordinates of lateral occipital cortex in humans. J Neurosurg. 2004 Jul;101(1):114–8.

384. Sayres R, Grill-Spector K. Relating retinotopic and object-selective responses in human lateral occipital cortex. J Neurophysiol. 2008 Jul;100(1):249–67.

385. Schüz A, Palm G. Density of neurons and synapses in the cerebral cortex of the mouse. J Comp Neurol. 1989 Aug 22;286(4):442–55.

386. Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, et al. Frequencies Contributing to Functional Connectivity in the Cerebral Cortex in “Resting-state” Data. AJNR Am J Neuroradiol. 2001 Aug 1;22(7):1326–33.

387. Kriegeskorte N, Cusack R, Bandettini P. How does an fMRI voxel sample the neuronal activity pattern: compact-kernel or complex spatiotemporal filter? Neuroimage. Elsevier Inc.; 2010 Feb 1;49(3):1965–76.

388. Shmuel A, Chaimow D, Raddatz G, Ugurbil K, Yacoub E. Mechanisms underlying decoding at 7 T: ocular dominance columns, broad structures, and macroscopic blood vessels in V1 convey information on the stimulated eye. Neuroimage. 2010 Feb 1;49(3):1957–64.

389. Swisher JD, Gatenby JC, Gore JC, Wolfe BA, Moon C-H, Kim S-G, et al. Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. J Neurosci. 2010 Jan 6;30(1):325–30.

390. Gardner JL. Is cortical vasculature functionally organized? Neuroimage. 2010;Vol 49(3):1953–6.

391. Op de Beeck HP. Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? Neuroimage. Elsevier Inc.; 2010 Mar

1;49(3):1943–8.

392. Hansen LK, Larsen J, Nielsen FA, Strother SC, Rostrup E, Savoy R, et al. Generalizable patterns in neuroimaging: how many principal components? Neuroimage. 1999 May;9(5):534–44.

393. Mitchell JF, Stoner GR, Reynolds JH. Object-based attention determines dominance in binocular rivalry. Nature. 2004 May 27;429(6990):410–3.

394. Tootell RB, Hadjikhani NK, Mendola JD, Marrett S, Dale a M. From retinotopy to recognition: fMRI in human visual cortex. Trends Cogn Sci. 1998 May 1;2(5):174–83.

395. Brefczynski JA, DeYoe EA. A physiological correlate of the “spotlight” of visual attention. Nat Neurosci. 1999 Apr;2(4):370–4.

396. Gandhi SP, Heeger DJ, Boynton GM. Spatial attention affects brain activity in human primary visual cortex. Proc Natl Acad Sci. 1999 Mar 16;96(6):3314–9.

397. Somers DC, Dale AM, Seiffert AE, Tootell RB. Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. Proc Natl Acad Sci U S A. 1999 Feb 16;96(4):1663–8.

398. Silver MA, Ress D, Heeger DJ. Topographic maps of visual spatial attention in human parietal cortex. J Neurophysiol. 2005 Aug 1;94(2):1358–71.

399. Schluppeck D, Curtis CE, Glimcher PW, Heeger DJ. Sustained activity in topographic areas of human posterior parietal cortex during memory-guided saccades. J Neurosci. 2006 May 10;26(19):5098–108.

400. Nachev P, Kennard C, Husain M. Functional role of the supplementary and pre-supplementary motor areas. Nat Rev Neurosci. 2008 Nov;9(11):856–69.

401. Kraft A, Sommer WH, Schmidt S, Brandt S a. Dynamic upper and lower visual field preferences within the human dorsal frontoparietal attention network. Hum Brain Mapp. 2011 Jul;32(7):1036–49.

402. Danckert J, Goodale M. The ups and downs of visual perception. In: Johnson SH, editor. Cognitive neuroscience perspectives on the problem of intentional action. MIT Press, Cambridge, MA,; 2003. p. pp 29–64.

403. Previc FH. Functional specialization in the lower and upper visual fields in humans: Its ecological origins and neurophysiological implications. Behav

Brain Sci. Cambridge University Press; 1990 Sep 1;13(03):519–42.

404. Previc FH. The neuropsychology of 3-D space. Psychol Bull. 1998 Sep;124(2):123–64.

405. Sayres R, Grill-spector K, Gotts SJ, Milleville SC, Bellgowan PSF, Martin A, et al. Relating Retinotopic and Object-Selective Responses in Human Lateral Occipital Cortex Relating Retinotopic and Object-Selective Responses in Human Lateral Occipital Cortex. 2011;(May 2008):249–67.

406. Curcio CA, Sloan KR, Packer O, Hendrickson AE, Kalina RE. Distribution of cones in human and monkey retina: individual variability and radial asymmetry. Science. UNITED STATES; 1987 May 1;236(4801):579–82.

407. Schein S, de Monasterio F. Mapping of retinal and geniculate neurons onto striate cortex of macaque. J Neurosci. 1987 Apr 1;7(4):996–1009.

408. Galletti C, Fattori P, Gamberini M, Kutz DF. The cortical visual area V6: brain location and visual topography. Eur J Neurosci. 1999 Nov;11(11):3922–36.

409. Maunsell JH, Van Essen DC. Topographic organization of the middle temporal visual area in the macaque monkey: representational biases and the relationship to callosal connections and myeloarchitectonic boundaries. J Comp Neurol. 1987 Dec 22;266(4):535–55.

410. Carlson T, Hogendoorn H, Fonteijn H, Verstraten F a J. Spatial coding and invariance in object-selective cortex. Cortex. Elsevier Srl; 2011 Jan;47(1):14–22.

411. Laar B Van De, Hayrettin G, Bos DP, Nijboer F, Nijholt A. Towards Practical Brain-Computer Interfaces. 2013;(March 2016).

412. de Haas B, Schwarzkopf DS, Urner M, Rees G. Auditory modulation of visual stimulus encoding in human retinotopic cortex. Neuroimage. 2013 Apr 15;70:258–67.

413. Schwarzkopf DS, Sterzer P, Rees G. Decoding of coherent but not incoherent motion signals in early dorsal visual cortex. Neuroimage. Elsevier Inc.; 2011 May 15;56(2):688–98.

414. Erez Y, Duncan J. Discrimination of Visual Categories Based on Behavioral Relevance in Widespread Regions of Frontoparietal Cortex. J Neurosci.

2015;35(36):12383–93.

415. Macevoy SP, Epstein R a. Decoding the representation of multiple simultaneous objects in human occipitotemporal cortex. Curr Biol. Elsevier Ltd; 2009 Jun 9;19(11):943–7.

416. Carlson TA. High temporal resolution decoding of object position and category. 2011;11:1–17.

417. Kravitz DJ, Kriegeskorte N, Baker CI. High-Level Visual Object Representations Are Constrained by Position. 2010;

418. Carlson T a. Orientation decoding in human visual cortex: new insights from an unbiased perspective. J Neurosci. 2014 Jun 11;34(24):8373–83.

419. Singh AK, Phillips S. Hierarchical control of false discovery rate for phase locking measures of EEG synchrony. Neuroimage. 2010 Mar;50(1):40–7.

420. van Gerven M, Bahramisharif A, Heskes T, Jensen O. Selecting features for BCI control based on a covert spatial attention paradigm. Neural Netw. Elsevier Ltd; 2009 Nov;22(9):1271–7.

421. Birbaumer N, Murguialday AR, Cohen L. Brain-computer interface in paralysis. Curr Opin Neurol. 2008 Dec;21(6):634–8.

422. Kastner S, Ungerleider LG. Mechanisms of visual attention in the human cortex. Annu Rev Neurosci. 2000 Jan;23:315–41.

423. Serences JT, Yantis S. Selective visual attention and perceptual coherence. Trends Cogn Sci. 2006 Jan;10(1):38–45.

424. Reynolds JH, Chelazzi L. Attentional modulation of visual processing. Annu Rev Neurosci. 2004;27:611–47.

425. Silver MA, Kastner S. Topographic maps in human frontal and parietal cortex. 2009;(September).

426. Baluch F, Itti L. Mechanisms of top-down attention. Trends Neurosci. Elsevier Ltd; 2011 Apr;34(4):210–24.

427. Hamker FH. Modeling feature-based attention as an active top-down inference process. Biosystems. Jan;86(1-3):91–9.

428. Moore T, Armstrong KM. Selective gating of visual signals by microstimulation

of frontal cortex. Nature. 2003 Jan 23;421(6921):370–3.

429. Treue S, Martínez Trujillo JC. Feature-based attention influences motion processing gain in macaque visual cortex. Nature. 1999 Jun 10;399(6736):575–9.

430. Maunsell JHR, Treue S. Feature-based attention in visual cortex. Trends Neurosci. 2006 Jun;29(6):317–22.

431. Hung CP, Kreiman G, Poggio T, DiCarlo JJ. Fast readout of object identity from macaque inferior temporal cortex. Science. 2005 Nov 4;310(5749):863–6.

432. Vansteensel MJ, Hermes D, Aarnoutse EJ, Bleichner MG, Schalk G, van Rijen PC, et al. Brain-computer interfacing based on cognitive control. Ann Neurol. 2010 Jun;67(6):809–16.

433. Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, et al. Frequencies contributing to functional connectivity in the cerebral cortex in “resting-state” data. AJNR Am J Neuroradiol. 2001 Aug;22(7):1326–33.

434. Wandell B a, Dumoulin SO, Brewer A a. Visual field maps in human cortex. Neuron. 2007 Oct 25;56(2):366–83.

435. Maunsell JHR, Cook EP. The role of attention in visual processing. Philos Trans R Soc Lond B Biol Sci. 2002 Aug 29;357(1424):1063–72.

436. LaConte SM, Peltier SJ, Hu XP. Real-time fMRI using brain-state classification. Hum Brain Mapp. 2007 Oct;28(10):1033–44.

437. Friedrich EVC, Wood G, Scherer R, Neuper C. Mind over brain, brain over mind: cognitive causes and consequences of controlling brain activity. Front Hum Neurosci. 2014 Jan;8(May):348.

438. Kim JG, Biederman I, Juan C-H. The benefit of object interactions arises in the lateral occipital cortex independent of attentional modulation from the intraparietal sulcus: a transcranial magnetic stimulation study. J Neurosci. 2011 Jun 1;31(22):8320–4.

439. Kim JG, Kastner S. Attention flexibly alters tuning for object categories. Trends Cogn Sci. Elsevier Ltd; 2013 Aug;17(8):368–70.

440. Halgren E, Dale a M, Sereno MI, Tootell RB, Marinkovic K, Rosen BR. Location of human face-selective cortex with respect to retinotopic areas. Hum

Brain Mapp. 1999 Jan;7(1):29–37.

441. Slotnick SD, White RC. The fusiform face area responds equivalently to faces and abstract shapes in the left and central visual fields. Neuroimage. Elsevier Inc.; 2013 Dec;83:408–17.

442. Reddy L, Moradi F, Koch C. Top-down biases win against focal attention in the fusiform face area. Neuroimage. 2007 Dec;38(4):730–9.

443. Yi D-J, Kelley T a, Marois R, Chun MM. Attentional modulation of repetition attenuation is anatomically dissociable for scenes and faces. Brain Res. 2006 Mar 29;1080(1):53–62.

444. Esterman M, Chiu Y-C, Tamber-Rosenau BJ, Yantis S. Decoding cognitive control in human parietal cortex. Proc Natl Acad Sci U S A. 2009 Oct 20;106(42):17974–9.

445. Koenigs M, Barbey AK, Postle BR, Grafman J. Superior parietal cortex is critical for the manipulation of information in working memory. J Neurosci. 2009 Nov 25;29(47):14980–6.

446. Yantis S, Schwarzbach J, Serences JT, Carlson RL, Steinmetz MA, Pekar JJ, et al. Transient neural activity in human parietal cortex during spatial attention shifts. Nat Neurosci. 2002 Oct;5(10):995–1002.

447. Yantis S, Serences JT. Cortical mechanisms of space-based and object-based attentional control. Curr Opin Neurobiol. 2003 Apr;13(2):187–93.

448. Liu T, Slotnick SD, Serences JT, Yantis S. Cortical mechanisms of feature-based attentional control. Cereb Cortex. 2003 Dec;13(12):1334–43.

449. Giesbrecht B, Woldorff MG, Song AW, Mangun GR. Neural mechanisms of top-down control during spatial and feature attention. Neuroimage. 2003 Jul;19(3):496–512.

450. Treisman A. Feature binding, attention and object perception. Philos Trans R Soc London - Ser B Biol Sci. 1998;353(1373):1295–306.

451. Bouvier S, Treisman A. Visual feature binding requires reentry. Psychol Sci. 2010;21(2):200–4.

452. Shomstein S, Behrmann M. Cortical systems mediating visual attention to both objects and spatial locations. Proc Natl Acad Sci U S A. 2006 Jul

25;103(30):11387–92.

453. Shomstein S, Lee J, Behrmann M. Top-down and bottom-up attentional guidance: investigating the role of the dorsal and ventral parietal cortices. Exp Brain Res. 2010 Oct;206(2):197–208.

454. Shomstein S. Cognitive functions of the posterior parietal cortex: top-down and bottom-up attentional control. Front Integr Neurosci. 2012 Jan;6(July):38.

455. Macevoy SP, Yang Z. Joint neuronal tuning for object form and position in the human lateral occipital complex. Neuroimage. Elsevier Inc.; 2012 Aug 24;63(4):1901–8.

456. Kravitz DJ, Kriegeskorte N, Baker CI. High-level visual object representations are constrained by position. Cereb Cortex. 2010 Dec;20(12):2916–25.

457. MacEvoy SP, Epstein R a. Constructing scenes from objects in human occipitotemporal cortex. Nat Neurosci. Nature Publishing Group; 2011 Oct;14(10):1323–9.

458. Kravitz DJ, Behrmann M. Space-, object-, and feature-based attention interact to organize visual scenes. Atten Percept Psychophys. 2011 Nov;73(8):2434–47.

459. Coull JT, Nobre a C. Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. J Neurosci. 1998 Sep 15;18(18):7426–35.

460. Assmus A, Marshall JC, Ritzl A, Noth J, Zilles K, Fink GR. Left inferior parietal cortex integrates time and space during collision judgments. Neuroimage. 2003 Nov;20:S82–8.

461. Assmus a, Marshall JC, Noth J, Zilles K, Fink GR. Difficulty of perceptual spatiotemporal integration modulates the neural activity of left inferior parietal cortex. Neuroscience. 2005 Jan;132(4):923–7.

462. Lewis P a, Miall RC. Distinct systems for automatic and cognitively controlled time measurement: evidence from neuroimaging. Curr Opin Neurobiol. 2003 Apr;13(2):250–5.

463. Lewis, Penelope a, Walsh V. Time Perception: Components of the Brain's clock. Curr Biol. 2005 May 24;15(10):R387–9.

464. Battelli L, Walsh V, Pascual-Leone A, Cavanagh P. The “when” parietal pathway explored by lesion studies. Curr Opin Neurobiol. 2008 Apr;18(2):120–6.

465. Bueti D, Bahrami B, Walsh V. Sensory and association cortex in time perception. J Cogn Neurosci. 2008 Jun;20(6):1054–62.

466. Kelly S., Lalor E, Reilly R., Foxe J. Independent Brain Computer Interface Control using Visual Spatial. Proc 2nd Int IEEE EMBS. 2005;

467. Kübler A, Kotchoubey B, Perelmouter J, Schauer M. The thought translation device : a neurophysiological approach to communication in total motor paralysis. 1999;223–32.

468. Kübler A, Neumann N, Kaiser J, Kotchoubey B, Hinterberger T, Birbaumer NP. Brain-computer communication: Self-regulation of slow cortical potentials for verbal communication. Arch Phys Med Rehabil. 2001 Nov;82(11):1533–9.

469. Kübler A, Neumann N, Wilhelm B, Hinterberger T, Birbaumer N. Predictability of Brain-Computer Communication. 2004;18:121–9.

470. Niazi AM, van den Broek PLC, Klanke S, Barth M, Poel M, Desain P, et al. Online decoding of object-based attention using real-time fMRI. Eur J Neurosci. 2014 Jan;39(2):319–29.

471. Barker-Collo SL. Within session practice effects on the PASAT in clients with multiple sclerosis. Arch Clin Neuropsychol. 2005 Mar;20(2):145–52.

472. Boksem M a S, Meijman TF, Lorist MM. Effects of mental fatigue on attention: an ERP study. Brain Res Cogn Brain Res. 2005 Sep;25(1):107–16.

473. Cook DB, O’Connor PJ, Lange G, Steffener J. Functional neuroimaging correlates of mental fatigue induced by cognition among chronic fatigue syndrome patients and controls. Neuroimage. 2007 May 15;36(1):108–22.

474. Faber LG, Maurits NM, Lorist MM. Mental fatigue affects visual selective attention. PLoS One. 2012 Jan;7(10):e48073.

475. Corbetta M, Akbudak E, Conturo TE, Snyder AZ, Ollinger JM, Drury HA, et al. A common network of functional areas for attention and eye movements. Neuron. 1998 Oct;21(4):761–73.

476. Gunduz A, Brunner P, Daitch A, Leuthardt EC, Ritaccio AL, Pesaran B, et al.

Decoding covert spatial attention using electrocorticographic (ECoG) signals in humans. Neuroimage. Elsevier Inc.; 2012 May 1;60(4):2285–93.

477. Treder MS, Bahramisharif A, Schmidt NM, van Gerven MAJ, Blankertz B. Brain-computer interfacing using modulations of alpha activity induced by covert shifts of attention. J Neuroeng Rehabil. 2011 Jan;8:24.

478. Peelen M V, Kastner S. Attention in the real world: toward understanding its neural basis. Trends Cogn Sci. Elsevier Ltd; 2014 May;18(5):242–50.

479. Suzuki S, Grabowecky M. Long-term speeding in perceptual switches mediated by attention-dependent plasticity in cortical visual processing. Neuron. 2007 Nov 21;56(4):741–53.

480. Folk C, Remington RW, Johnston CJ. Involuntary Covert Orienting Is Contigent on Attentional Control Settings. J Exp Psychol. 1992;18(4):1030–44.

481. Seidl-Rathkopf KN, Turk-Browne NB, Kastner S. Automatic guidance of attention during real-world visual search. Atten Percept Psychophys. 2015 Aug;77(6):1881–95.

482. Peelen M V, Kastner S. A neural basis for real-world visual search in human occipitotemporal cortex. Proc Natl Acad Sci U S A. 2011 Jul 19;108(29):12125–30.

483. Hanakawa T, Immisch I, Toma K, Dimyan M a, Van Gelderen P, Hallett M. Functional properties of brain areas associated with motor execution and imagery. J Neurophysiol. 2003 Feb;89(2):989–1002.

484. Grezes J, Decety J. Does visual perception of object afford action ? Evidence from a neuroimaging study. Neuropsychologia. 2002;40:212–22.

485. Pearson J, Clifford CWG, Tong F. The functional impact of mental imagery on conscious perception. Curr Biol. 2008 Jul 8;18(13):982–6.

486. Rademaker RL, Pearson J. Training Visual Imagery: Improvements of Metacognition, but not Imagery Strength. Front Psychol. 2012 Jan;3(July):224.

487. Foerster Á, Rocha S, Wiesiolek C, Chagas AP, Machado G, Silva E, et al. Site-specific effects of mental practice combined with transcranial direct current stimulation on motor learning. Eur J Neurosci. 2013 Mar 19;37(5):786–94.

488. Decety J, Perani D, Jeannerod M, Bettinardi V, Tadary B, Woods R, et al. Mapping Motor Representations with positron emission tomography. Nature. Nature Publishing Group; 1994;371(13):600–2.

489. Munzert J, Lorey B, Zentgraf K. Cognitive motor processes: the role of motor imagery in the study of motor representations. Brain Res Rev. Elsevier B.V.; 2009 May;60(2):306–26.

490. Levelt WJM. The alternation process in binocular rivalry. Br J Psychol. 1966;57(3):225–38.

491. Lin Z, He S. Seeing the invisible: the scope and limits of unconscious processing in binocular rivalry. Prog Neurobiol. 2009 Apr;87(4):195–211.

492. Carmel D, Arcaro M, Kastner S, Hasson U. How to create and use binocular rivalry. J Vis Exp. 2010 Jan;(45):1–10.

493. Blake R, Logothetis NK. Visual competition. Nat Rev Neurosci. 2002 Jan;3(1):13–21.

494. Wheatstone C. On some remarkable and hitherto unobserved phenomena of binocular vision. Optom Wkly. 1962;53:2311–5.

495. Fox R, Herrmann J. Stochastic properties of. 1967;2(9):432–6.

496. Tong F, Meng M, Blake R. Neural bases of binocular rivalry. Trends Cogn Sci. 2006 Nov;10(11):502–11.

497. Dayan P. A hierarchical model of binocular rivalry. Neural Comput. 1998 Jul 1;10(5):1119–35.

498. Parker A, Alais D. A bias for looming stimuli to predominate in binocular rivalry. Vision Res. 2007 Sep;47(20):2661–74.

499. Kanai R, Carmel D, Bahrami B, Rees G. Structural and functional fractionation of right superior parietal cortex in bistable perception. Curr Biol. Elsevier; 2011 Feb 8;21(3):R106–7.

500. Kanai R, Bahrami B, Rees G. Human parietal cortex structure predicts individual differences in perceptual rivalry. Curr Biol. Elsevier Ltd; 2010 Sep 28;20(18):1626–30.

501. Deichmann R, Schwarzbauer C, Turner R. Optimisation of the 3D MDEFT

sequence for anatomical brain imaging: technical implications at 1.5 and 3 T. Neuroimage. 2004 Feb;21(2):757–67.

502. Stephan KE, Penny WD, Moran RJ, den Ouden HEM, Daunizeau J, Friston KJ. Ten simple rules for dynamic causal modeling. Neuroimage. Elsevier Inc.; 2010 Feb 15;49(4):3099–109.

503. Friston KJ, Harrison L, Penny W. Dynamic causal modelling. Neuroimage. 2003 Aug;19(4):1273–302.

504. Marreiros a C, Kiebel SJ, Friston KJ. Dynamic causal modelling for fMRI: a two-state model. Neuroimage. 2008 Jan 1;39(1):269–78.

505. Stephan KE, Penny WD, Moran RJ, den Ouden HEM, Daunizeau J, Friston KJ. Ten simple rules for dynamic causal modeling. Neuroimage. 2010;49(4):3099–109.

506. Friston K, Penny W. Post hoc Bayesian model selection. Neuroimage. Elsevier Inc.; 2011 Jun 15;56(4):2089–99.

507. Friston KJ, Frith CD, Frackowiak RSJ, Turner R. Characterizing Dynamic Brain Responses with fMRI: A Multivariate Approach. Neuroimage. 1995 Jun;2(2):166–72.

508. Friston KJ, Kahan J, Biswal B, Razi A. A DCM for resting state fMRI. Neuroimage. 2014 Jul 1;94:396–407.

509. Shibata K, Watanabe T, Sasaki Y, Kawato M. Perceptual Learning Incepted by Decoded fMRI Neurofeedback Without Stimulus Presentation. Science (80-). American Association for the Advancement of Science; 2011 Dec 9;334(6061):1413–5.

510. Kim S, Rauss K, Yu T, Dax J, Muehlbauer F, Soekadar SR, et al. Training the 6th sense: making subliminal stimuli perceivable using real-time fMRI. 2011.

511. Sulzer J, Haller S, Scharnowski F, Weiskopf N, Birbaumer N, Blefari ML, et al. Real-time fMRI neurofeedback: progress and challenges. Neuroimage. Elsevier Inc.; 2013 Aug 1;76:386–99.

512. Berman BD, Horovitz SG, Venkataraman G, Hallett M. NeuroImage Self-modulation of primary motor cortex activity with motor and motor imagery tasks using real-time fMRI-based neurofeedback. Neuroimage. Elsevier Inc.;

2012;59(2):917–25.

513. Ninaus M, Kober SE, Witte M, Koschutnig K, Stangl M, Neuper C, et al. Neural substrates of cognitive control under the belief of getting neurofeedback training. Front Hum Neurosci. 2013 Jan;7:914.

514. Kober SE, Witte M, Ninaus M, Neuper C, Wood G. Learning to modulate one's own brain activity: the effect of spontaneous mental strategies. Front Hum Neurosci. 2013;7(October):1–12.

515. Bagdasaryan J, Quyen ML Van. Experiencing your brain: neurofeedback as a new bridge between neuroscience and phenomenology. Front Hum Neurosci. 2013 Jan;7(October):680.

516. Kühn S, Brass M, Haggard P. Feeling in control: Neural correlates of experience of agency. Cortex. 2012;49(7):1935–42.

517. Iseki K, Hanakawa T, Shinozaki J, Nankaku M, Fukuyama H. Neural mechanisms involved in mental imagery and observation of gait. Neuroimage. 2008 Jul 1;41(3):1021–31.

518. Carmel DP. Top-down control of visual attention and awareness : Cognitive and neural mechanisms. 2006;(September).

519. Friston K. Causal modelling and brain connectivity in functional magnetic resonance imaging. PLoS Biol. Public Library of Science; 2009 Feb 17;7(2):e33.

520. Felleman DJ, Van Essen DC. Distributed Hierarchical Processing in the Primate Cerebral Cortex. Cereb Cortex. 1991 Jan 1;1(1):1–47.

521. Rosa MJ, Friston K, Penny W. Post-hoc selection of dynamic causal models. J Neurosci Methods. Elsevier B.V.; 2012 Jun 30;208(1):66–78.

522. Frässle S, Sommer J, Jansen A, Naber M, Einhäuser W. Binocular rivalry: frontal activity relates to introspection and action but not to perception. J Neurosci. 2014 Jan 29;34(5):1738–47.

523. Theodoni P, Panagiotaropoulos TI, Kapoor V, Logothetis NK, Deco G. Cortical microcircuit dynamics mediating binocular rivalry: the role of adaptation in inhibition. Front Hum Neurosci. 2011 Jan;5(November):145.

524. Zaretskaya N, Thielscher A, Logothetis NK, Bartels A. Disrupting parietal

function prolongs dominance durations in binocular rivalry. Curr Biol. 2010 Dec 7;20(23):2106–11.

525. Marx S. Reward modulates perception in binocular rivalry. 2015;15:1–13.

526. Slotnick SD, Yantis S. Common neural substrates for the control and effects of visual attention and perceptual bistability. Brain Res Cogn Brain Res. 2005 Jun;24(1):97–108.

527. Hohwy J, Roepstorff A, Friston K. Predictive coding explains binocular rivalry: an epistemological review. Cognition. 2008 Sep;108(3):687–701.

528. Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L. Stimulus predictability reduces responses in primary visual cortex. J Neurosci. 2010 Feb 24;30(8):2960–6.

529. Denison RN, Piazza E a, Silver M a. Predictive Context Influences Perceptual Selection during Binocular Rivalry. Front Hum Neurosci. 2011 Jan;5(December):166.

530. Pelekanos V, Roumani D, Moutoussis K. The effects of categorical and linguistic adaptation on binocular rivalry initial dominance. Front Hum Neurosci. 2011 Jan;5(January):187.

531. Sagi Y, Tavor I, Hofstetter S, Tzur-Moryosef S, Blumenfeld-Katzir T, Assaf Y. Learning in the Fast Lane: New Insights into Neuroplasticity. Neuron. Elsevier Inc.; 2012;73(6):1195–203.

532. Johansen-Berg H, Baptista CS, Thomas AG. Human structural plasticity at record speed. Neuron. Elsevier Inc.; 2012 Mar 22;73(6):1058–60.

533. Hilti CC, Jann K, Heinemann D, Federspiel A, Dierks T, Seifritz E, et al. Evidence for a cognitive control network for goal-directed attention in simple sustained attention. Brain Cogn. 2013 Mar;81(2):193–202.

534. Bonini F, Burle B, Liégeois-Chauvel C, Régis J, Chauvel P, Vidal F. Action monitoring and medial frontal cortex: leading role of supplementary motor area. Science. 2014 Feb 21;343(6173):888–91.

535. Pearson J, Naselaris T, Holmes EA, Kosslyn SM. Mental Imagery : Functional Mechanisms and Clinical Applications. 2015;19(10):590–602.

536. Bacher D, Jarosiewicz B, Masse NY, Stavisky SD, Simeral JD, Newell K, et al.

Neural Point-and-Click Communication by a Person With Incomplete Locked-In Syndrome. Neurorehabil Neural Repair. 2015 Jun;29(5):462–71.

537. Brascamp JW, Ee R Van, Noest J, Jacobs RHAH, Berg AV Van Den. The time course of binocular rivalry reveals a fundamental role of noise. 2006;1244–56.

538. Klink PC, van Ee R, van Wezel RJ a. General validity of Levelt's propositions reveals common computational mechanisms for visual rivalry. PLoS One. 2008 Jan;3(10):e3473.

539. Dieter KC, Tadin D. Understanding attentional modulation of binocular rivalry: a framework based on biased competition. Front Hum Neurosci. 2011 Jan;5(December):155.

540. Shibata K, Sagi D, Watanabe T. Two-stage model in perceptual learning: toward a unified theory. Ann N Y Acad Sci. 2014 May;1316:18–28.

541. Koush Y, Rosa MJ, Robineau F, Heinen K, Rieger S, Weiskopf N, et al. Connectivity-based neurofeedback: Dynamic causal modeling for real-time fMRI. Neuroimage. Elsevier Inc.; 2013 May 10;81:422–30.

542. Ghaziri J, Tucholka A, Larue V, Blanchette-Sylvestre M, Reyburn G, Gilbert G, et al. Neurofeedback training induces changes in white and gray matter. Clin EEG Neurosci. 2013 Oct;44(4):265–72.

543. Ishai A, Sagi D. Common mechanisms of visual imagery and perception. Science (80-). 1995 Jun 23;268(5218):1772–4.

544. Mohr HM, Linder NS, Dennis H, Sireteanu R. Orientation-Specific Aftereffects to Mentally Generated Lines. Perception. 2011 Mar 1;40(3):272–90.

545. Winawer J, Huk AC, Boroditsky L. A motion aftereffect from visual imagery of motion. Cognition. 2010 Feb;114(2):276–84.

546. Pearson J, Rademaker RL, Tong F. Evaluating the mind's eye: the metacognition of visual imagery. Psychol Sci. 2011 Dec;22(12):1535–42.

547. Orban G a, Zhu Q, Vanduffel W. The transition in the ventral stream from feature to real-world entity representations. Front Psychol. 2014 Jan;5(July):695.

548. Miranda RA, Casebeer WD, Hein AM, Judy JW, Krotkov EP, Laabs TL, et al. DARPA-funded efforts in the development of novel brain – computer interface

technologies. 2015;244:52–67.

549. Johansen-Berg H. High field MRI for study of adaptive brain plasticity. Wellcome Trust Brain Meetings. 2015.

550. Yousry TA, Schmid UD, Alkadhi H, Schmidt D, Peraud A, Buettner A, et al. Localization of the motor hand area to a knob on the precentral gyrus A new landmark. 1997;141–57.

551. Hiremath S V, Chen W, Wang W, Foldes S, Yang Y, Tyler-Kabara EC, et al. Brain computer interface learning for systems based on electrocorticography and intracortical microelectrode arrays. Front Integr Neurosci. 2015 Jan;9(June):40.

552. DiGiovanna J, Mahmoudi B, Fortes J, Principe JC, Sanchez JC. Coadaptive brain-machine interface via reinforcement learning. IEEE Trans Biomed Eng. 2009 Jan;56(1):54–64.

553. Sampaio-Baptista C, Khrapitchev AA, Foxley S, Schlagheck T, Scholz J, Jbabdi S, et al. Motor skill learning induces changes in white matter microstructure and myelination. J Neurosci. 2013 Dec 11;33(50):19499–503.

554. Takeuchi H, Sekiguchi A, Taki Y, Yokoyama S, Yomogida Y, Komuro N, et al. Training of working memory impacts structural connectivity. J Neurosci. 2010 Mar 3;30(9):3297–303.

555. Yamada K, Miyawaki Y, Kamitani Y. Inter-subject neural code converter for visual image representation. Neuroimage. Elsevier Inc.; 2015 Jun;113:289–97.

556. Rao RPN, Stocco A, Bryan M, Sarma D, Youngquist TM, Wu J, et al. A direct brain-to-brain interface in humans. PLoS One. 2014 Jan;9(11):e111332.

557. Grau C, Ginhoux R, Riera A, Nguyen TL, Chauvat H, Berg M, et al. Conscious Brain-to-Brain Communication in Humans Using Non-Invasive Technologies. PLoS One. 2014 Jan;9(8):e105225.

558. Pais-Vieira M, Lebedev M, Kunicki C, Wang J, Nicolelis MAL. A Brain-to-Brain Interface for Real-Time Sharing of Sensorimotor Information. Sci Rep. Nature Publishing Group; 2013 Feb 28;3:1319.

559. Blabe CH, Gilja V, Chestek C a, Shenoy K V, Anderson KD, Henderson JM. Assessment of brain-machine interfaces from the perspective of people with

paralysis. J Neural Eng. IOP Publishing; 2015 Jul 14;12(4):043002.

560. Popescu F, Fazli S, Badower Y, Blankertz B, Müller K-R. Single Trial Classification of Motor Imagination Using 6 Dry EEG Electrodes. Miall C, editor. PLoS One. Public Library of Science; 2007 Jul 25;2(7):e637.

561. Ruffini G, Dunne S, Farres E, Watts PCP, Mendoza E, Silva SRP, et al. ENOBIO - First Tests of a Dry Electrophysiology Electrode using Carbon Nanotubes. 2006 International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE; 2006. p. 1826–9.

562. Ruffini G, Dunne S, Farres E, Cester I, Watts PCP, Silva SRP, et al. ENOBIO dry electrophysiology electrode; first human trial plus wireless electrode system. 2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE; 2007. p. 6689–93.

563. Mihajlovic V, Grundlehner B, Vullers R, Penders J. Wearable, wireless EEG solutions in daily life applications: what are we missing? IEEE J Biomed Heal informatics. IEEE; 2015 Jan 1;19(1):6–21.

564. Millán JDR, Rupp R, Müller-Putz GR, Murray-Smith R, Giugliemma C, Tangermann M, et al. Combining Brain-Computer Interfaces and Assistive Technologies: State-of-the-Art and Challenges. Front Neurosci. 2010 Jan;4(September):1–15.

565. Ibarra-Orozco R, Gonzalez-Mendoza M, Hernandez-Gress N, Diederichs F, Kortelainen J. Towards a Ready-to-Use Drivers' Vigilance Monitoring System. 2008 International Conference on Computational Intelligence for Modelling Control & Automation. IEEE; 2008. p. 802–7.

566. Barrese JC, Rao N, Paroo K, Triebwasser C, Vargas-Irwin C, Franquemont L, et al. Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. J Neural Eng. IOP Publishing; 2013 Dec 1;10(6):066014.

567. Lu CW, Patil PG, Chestek C a. Current challenges to the clinical translation of brain machine interface technology. 1st ed. International review of neurobiology. Elsevier Inc.; 2012. 137-60 p.

568. Linden D. Brain Control: Developments in Therapy and Implications for

Society. Palgrave Macmillan; 2014. 208 p.

569. Awan NR, Lozano A, Hamani C. Deep brain stimulation: current and future perspectives. Neurosurg Focus. 2009 Jul;27(1):E2.

570. Hariz M, Blomstedt P, Zrinzo L. Future of brain stimulation: new targets, new indications, new technology. Mov Disord. 2013 Nov;28(13):1784–92.

571. Matharu MS, Zrinzo L. Deep brain stimulation in cluster headache. 2011;11(4):473–5.

572. Shine JM, Keogh R, Callaghan CO, Muller AJ, Lewis SJG, Pearson J. Imagine that : elevated sensory strength of mental imagery in individuals with Parkinson ' s disease and visual hallucinations. 2015;

573. Gratwicke J, Kahan J, Zrinzo L, Hariz M, Limousin P, Foltynie T, et al. The nucleus basalis of Meynert: A new target for deep brain stimulation in dementia? Neurosci Biobehav Rev. Elsevier Ltd; 2013 Dec;37(10):2676–88.

574. Pollok B, Kamp D, Butz M, Wojtecki L, Timmermann L, Südmeyer M, et al. Increased SMA-M1 coherence in Parkinson's disease - Pathophysiology or compensation? Exp Neurol. Elsevier Inc.; 2013 Sep;247:178–81.

575. Rektor I, Bočková M, Chrastina J, Rektorová I, Baláž M. The modulatory role of subthalamic nucleus in cognitive functions - a viewpoint. Clin Neurophysiol. 2015 Apr;126(4):653–8.

576. Nambu A. A new dynamic model of the cortico-basal ganglia loop. Prog Brain Res. 2004 Jan;143(03):461–6.

577. Priori A, Foffani G, Rossi L, Marceglia S. Adaptive deep brain stimulation (aDBS) controlled by local field potential oscillations. Exp Neurol. Elsevier B.V.; 2013 Jul;245:77–86.

578. Little S, Pogosyan A, Neal S, Zavala B, Zrinzo L, Hariz M, et al. Adaptive Deep Brain Stimulation in Advanced Parkinson Disease. 2013;1–9.

579. Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I. Invariant visual representation by single neurons in the human brain. Nature. Macmillian Magazines Ltd.; 2005 Jun 23;435(7045):1102–7.

580. Mesgarani N, Cheung C, Johnson K, Chang EF. Phonetic Feature Encoding in Human Superior Temporal Gyrus. Science (80-). 2014 Jan

30;343(6174):1006–10.

581. Leuthardt EC, Gaona C, Sharma M, Szrama N, Roland J, Freudenberg Z, et al. Using the electrocorticographic speech network to control a brain-computer interface in humans. J Neural Eng. 2011 Jun;8(3):036004.

582. Bouchard KE, Chang EF. N eural decoding of spoken vowels from human sensory - motor cortex with high - density electrocorticography. 2014;6782–5.

583. Wang W, Collinger JL, Degenhart AD, Tyler-Kabara EC, Schwartz AB, Moran DW, et al. An electrocorticographic brain interface in an individual with tetraplegia. PLoS One. 2013 Jan;8(2):e55344.

584. Lang PJ, Twentyman CT. Learning to Control Heart Rate: Effects of Varying Incentive and Criterion of Success on Task Performance. Psychophysiology. 1976 Sep;13(5):378–85.

585. Fried R, Fox MC, Carlton RM. Effect of diaphragmatic respiration with end-tidal CO2 biofeedback on respiration, EEG, and seizure frequency in idiopathic epilepsy. Ann N Y Acad Sci. UNITED STATES; 1990 Jan;602:67–96.

586. Wauquier A, McGrady A, Aloe L, Klausner T, Collins B. Changes in cerebral blood flow velocity associated with biofeedback-assisted relaxation treatment of migraine headaches are specific for the middle cerebral artery. Headache. UNITED STATES; 1995 Jun;35(6):358–62.

587. Duschek S, Schuepbach D, Doll A, Werner NS, Reyes del Paso G a. Self-regulation of cerebral blood flow by means of transcranial Doppler sonography biofeedback. Ann Behav Med. 2011 Apr;41(2):235–42.

588. Ciancarelli I, Tozzi-Ciancarelli MG, Spacca G, Di Massimo C, Carolei A. Relationship between biofeedback and oxidative stress in patients with chronic migraine. Cephalalgia. England; 2007 Oct;27(10):1136–41.

589. Critchley H. Volitional Control of Autonomic Arousal: A Functional Magnetic Resonance Study. Neuroimage. 2002 Aug;16(4):909–19.

590. Critchley HD, Harrison NA. Perspective Visceral Influences on Brain and Behavior. 2013;

591. Jones CL, Minati L, Nagai Y, Medford N, Harrison N a, Gray M, et al. Neuroanatomical substrates for the volitional regulation of heart rate. Front

Psychol. 2015 Jan;6(March):300.

592. Ekanayake J, Chari A, Craven C, Thompson SD, Shah SN, Patel NA, et al. Learning to control ICP. Fluids Barriers CNS. BioMed Central; 2015;12(Suppl 1):P12.