Visuospatial Attention based Brain-Computer Interfacing using Real-Time fMRI

Patrik Andersson

Visuospatial Attention based Brain-Computer Interfacing using Real-Time fMRI

PhD thesis, Utrecht University, the Netherlands.

ISBN: 978-90-8891-433-1

Print: Proefschriftmaken.nl

This book was typeset using LaTeX 2ε

VISUOSPATIAL ATTENTION BASED BRAIN-COMPUTER INTERFACING USING REAL-TIME fMRI

Brein-Computer Interactie gebaseerd op Visuospatiële Aandacht met behulp van Real-Time fMRI

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus prof.dr. G.J. van der Zwaan, ingevolge het
besluit van het college voor promoties in het openbaar te verdedigen
op woensdag 4 juli 2012 des middags te 4.15 uur

door

Jonas Patrik Andersson

geboren op 31 juli 1974
te Göteborg, Zweden

Promotoren: Prof. dr. ir. M.A. Viergever
Prof. dr. N.F. Ramsey

Co-promotor: Dr. J.P.W. Pluim

This thesis was partly accomplished with financial support from Philips Healthcare, Röntgen Stichting Utrecht and RoboRealm.

"I used to think that the human brain was
the most fascinating part of the body. Then I thought,
'what part of my body is telling me that?'"
Emo Philips

Contents

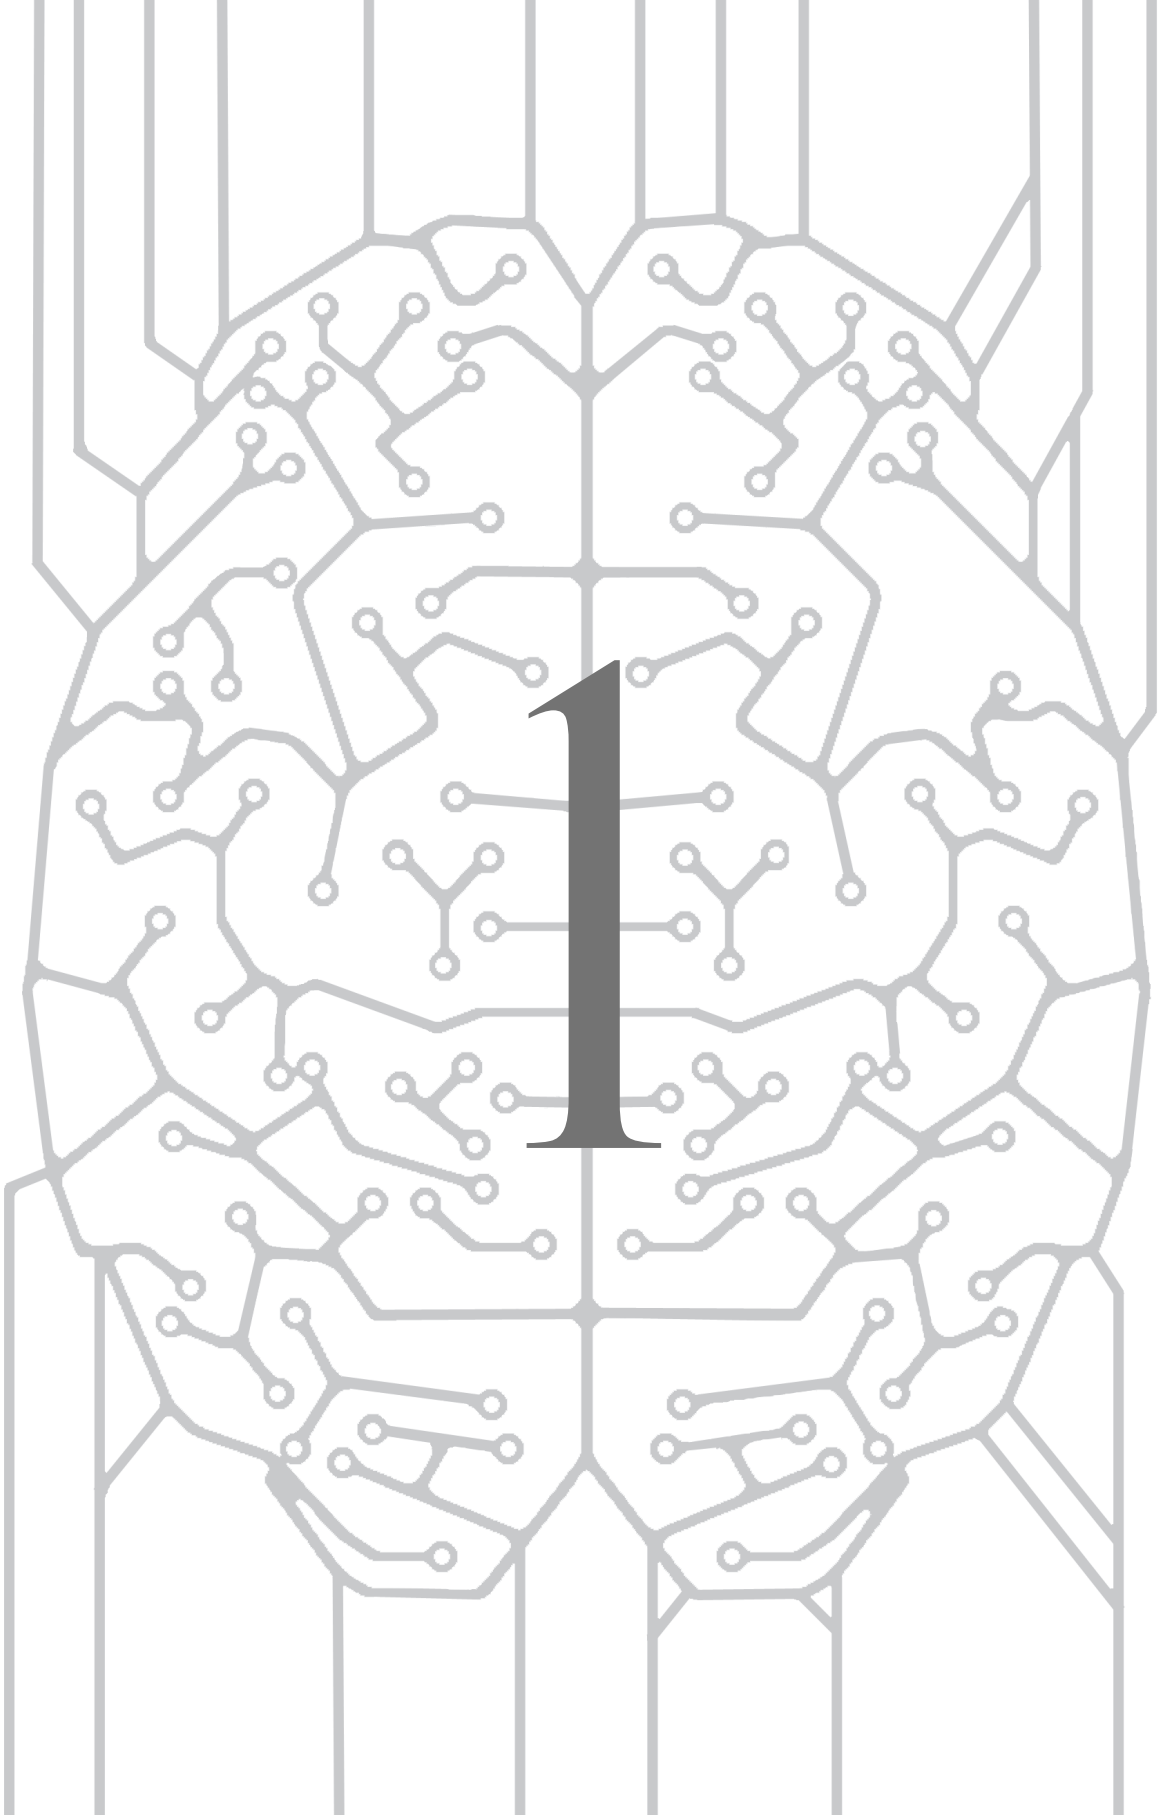
1

General introduction

1.1 Locked-in syndrome

"Thirty years ago a stroke left me in a coma. When I awoke I found myself completely paralyzed and unable to speak. I didn't know what paralysis was until I could move nothing but my eyes. I didn't know what loneliness was until I had to wait all night in the dark, in pain from head to foot, vainly hoping for someone to come with a teardrop of comfort. I didn't know what silence was until the only sound I could make was that of my own breath issuing from a hole drilled into my throat."

Julia Tavalaro - Look up for yes [1]

The quote above describes one of the most frightening scenarios one can imagine; to find oneself in a fully awake and conscious state of mind, but with no ability to communicate or move the body. This is the reality for people diagnosed with locked-in syndrome (LIS) [2]. LIS is a rare condition that can have several possible causes. Most common is an infarct or haemorrhage in the pons, located in the brainstem. Other causes can be neuronal damage via a tumor or trauma or a demyelinating disease such as Amyotrophic Lateral Sclerosis (ALS).

The quote above is from Julia Tavalaro who in 1966 woke up in a locked-in state after a seven months coma following a stroke. After she woke up, it took six years until it was noticed that she was aware. All this time, while her sensory systems functioned, her mind was imprisoned in a body she had no control over.

LIS has been classified into three different types [3] (Table 1.1), depending on the extent of motor impairment. The ability to communicate depends on which category a LIS patient belongs to. Since any available muscular control can potentially be used for communication, incomplete LIS patients are most likely to be able to maintain some form of communication.

After it was realized that Julia Tavalaro was aware and had some control over the eyes, she learned to communicate by responding with eye movements to someone pointing at letters on an alphabet board and eventually became a published poet and

Table 1.1: Categories of Locked-in Syndrome.

Classical LIS	Total immobility, with preserved vertical eye movements or blinking
Incomplete LIS	As classical but with remnants of voluntary movement other than vertical eye movement
Total LIS	Total immobility, including all eye movements

author. She is therefore a good example showing that a person can continue to be productive if given a means to communicate. Julia Tavalaro is also an example of a locked-in patient being misdiagnosed as being in a vegetative state. A patient in a persistent vegetative state may be awake but is without awareness. Studies have shown that the percentage of patients being misdiagnosed to be in a vegetative state can be as high as 40% [4]. Although this includes patients with any true level of awareness, it is likely that when the tools and techniques for diagnosis improve, the number of people that can benefit from assistive communication devices, such as Brain-Computer Interfaces, will grow.

1.2 Brain-Computer Interfaces

A Brain-Computer Interface (BCI) classifies neurophysiological signals into commands. By wilful regulation of local brain activity, a user can communicate intent without depending on any form of muscular control. The closed information loop in a BCI is illustrated in Figure 1.1. The cycle starts with the BCI user performing a mental task that represents a command. The brain activity is monitored using some neuroimaging technique, e.g. EEG or fMRI. After preprocessing of the raw data to remove or reduce acquisition artifacts, the relevant features can be extracted from the data and used for classifying the user's brain state. The result of the classification, hopefully reflecting the user's intent, is sent as a command to a device that updates accordingly. In the end, some form of feedback (e.g. visual, auditory or tactile) is given to the subject who can now take action based on the new circumstances.

With access to a functioning BCI, the quality of life for LIS patients could be drastically increased. For instance, control of the cursor on a computer screen would give access to most available software and the ability to select letters on a virtual keyboard would make it possible to communicate via writing. Also patients with less severe paralysis, or even healthy people, could benefit from BCI technologies. However, a BCI is only truly useful when it is equally, or more, efficient than the alternative tools available. With no residual motor function left to build another solution on, the first group to benefit from BCI is the group suffering from total LIS. For them, control of a simple binary switch would already be a great improvement. On the other hand, to help an amputee be able to control a prosthetic limb with many degrees of freedom, a simple switch is not enough.

Two of the most defining properties of a BCI are the modality used for detecting the brain activity and the mental task used for regulating it. The progress of BCI

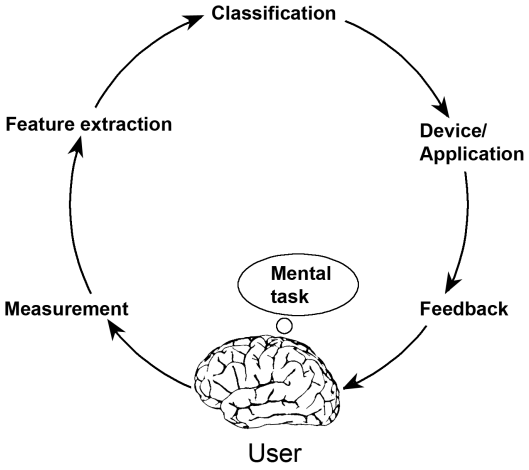

Figure 1.1: Illustration of the different modules in the BCI cycle. The cycle starts with the user regulating brain activity using a mental task and is closed by the feedback showing the effect.

therefore relies to a large part both on technological improvements and on a better understanding of the relationship between thoughts and the brain.

1.3 State of the art of BCI

Theoretically, any neuroimaging technique capable of detecting neural activity, referred to as functional neuroimaging, can be used for BCI. These techniques can be divided into two broad groups describing the type of measure used. The first group, including e.g. EEG and MEG, are based on electrophysiological measures of neural activity. These are direct measures of the electromagnetic field changes generated by current flows in firing neurons. The second group instead use an indirect hemodynamic measure and includes e.g. fMRI. The hemodynamic measures are based on the changes in metabolism due to neural activity. The increased energy consumption of firing neurons leads to changes in the blood flow, blood volume, oxygenation level etc that can be used for indirect measures of activity.

In practice, and to be clinically relevant, the BCI system needs to be portable and usable by patients at home. It should also be able to run without too much help from medical or technical personnel. These criteria are not fulfilled in the cases of

MEG and fMRI. There are however several descriptions of BCI systems built on both MEG [5, 6], and fMRI [7–11]. While these modalities can, as we will see, be very useful within BCI research, future clinical BCI systems used by patients will most likely be based on electrical signals.

1.3.1 Non-invasive BCI

Electroencephalography (EEG) uses electrodes placed on the scalp and is therefore completely non-invasive. EEG has so far been the predominant modality in BCI research. Besides being non-invasive, EEG systems are both portable and fairly easy to use.

To regulate the brain activity the BCI user needs to perform one, or multiple, mental control tasks. These control tasks must produce some specific pattern or spectral power change in the EEG signal that can be picked up and classified. The most common control strategies have been Slow Cortical Potentials (SCP) [12–14], the sensorimotor rhythm (SMR) [15], the P300 oddball response and Steady State Visually Evoked Potentials (SSVEP) [16–20].

The SCP BCI [21] has been shown to give rather low performances while needing long training periods and and it has therefore lost popularity.

The SMRs are EEG rhythms that change with movement or imagined movements. These rhythms can be modulated, and used for BCI control, by imagining moving parts of the body, for instance a hand or the tongue. In a paralyzed person the motor system has lost its functionality and is therefore "available" to be utilized by BCI and it could provide a very intuitive control of e.g. prosthetics.

Both P300 and SSVEP are evoked responses, i.e. they depend on an external stimulation to trigger them. The P300 response is evoked by a rare or surprising event and can be used for separating preferred and non-preferred stimuli. While a P300 response can be evoked via other sensory systems, the visual system is by far the most common for SSVEP BCI. A BCI based on P300 was first described in 1988 by Farwell and Donchin [22] and has since been extensively studied [23–26]. A (visual) P300 BCI has a computer screen displaying a number of choices that flash in a random sequence. When the user focuses on one of the choices, a P300 response will be evoked when this particular choice lights up.

SSVEPs are induced by flickering visual stimuli to which the user either moves the gaze, or focuses the attention. The frequency of the flickering is resonated in the EEG signals and in a SSVEP BCI different frequencies are linked to different choices. Some very encouraging results have been presented in studies using these control tasks

for EEG based BCI [21]. Despite this, the number of patients using BCI outside of research is quite limited. Although this could partly be blamed on the lack of commercial solutions, it is also a fact that the performance and reliability of the proposed systems are not good enough. It has been documented that around 20% of subjects testing these BCI systems can not attain a useable control [27–29]. In the case of motor based BCI there is also a possibility that the loss of voluntary motor control will lead to brain reorganization that in turn affects the BCI performance [21, 30]. Besides being intrinsically dependent on external visual stimulation, several studies have indicated that both P300 [31, 32] and SSVEP [32, 33] require eye movements to achieve acceptable control.

The main problem with EEG as technology is that the weak electrical potentials created by the firing neurons have to pass through several tissue layers, in particular the low conducting bone, before reaching the electrodes. The consequences are a substantial lost of signal power as well as a "smearing" effect resulting in a low spatial resolution. The tissues also act as a lowpass filter, making it hard to measure activity in the higher frequency bands.

As long as there are better solutions for patients with any residual muscle control, these will be used, and without a more reliable system, patients with progressive neurodegenerative deceases such as ALS will not choose to go into respiratory care and trust that a BCI can continue to execute their intentions when the muscles are no longer able to. If the limit of EEG BCI has more or less been reached, the only option is to switch to another modality. One alternative is to consider invasive methods in the form of implanted electrodes.

1.3.2 Invasive BCI

As the voltage field falls off with the square of the distance from the source, the signal power depends largely on how close to the neurons the electrodes are placed. With cortex-penetrating microelectrodes it is possible to monitor very small groups of, or even single, neurons directly at the source. These techniques have been applied for BCI research, mainly in animal studies [34–36] but there are also reports of successful BCI control in humans [37, 38]. The biggest issue with intracortical electrodes, besides the risk of infection, is the question of long term stability. When a foreign object is placed inside the gray matter there is a chance of fibrous tissue building up around it, which would weaken the signal over time. Further, measuring from such a small number of neurons also makes the system very sensitive to plastic changes in the brain.

In electrocorticography (ECoG) the recording electrodes are placed on the surface of the cortex. ECoG is less invasive than electrodes penetrating the cortex, and BCI systems based on this technology [39–43] are therefore sometimes referred to as partially-invasive. Because the ECoG electrodes have much less contact with the brain tissue compared to the more invasive microelectrodes, the technique is supposedly more stable over time [44].

Theoretically, any mental task processed within the sensory, motor and cognitive brain systems can be considered for BCI control. Since most BCI research has been conducted with EEG, the mental tasks evaluated have mainly been based on brain functions producing responses strong enough to be detected well from the scalp. There is currently a rapid increase in ECoG studies adding to the understanding of the basis of these signals. Parallel to this, there is a lot of progress in implant technology. Since invasive technologies make it possible to detect signal changes in brain regions not accessible with scalp electrodes, and at a much higher resolution, these advances can potentially open up for BCI control using additional strategies and brain functions [43, 45, 46]. A particular person might find some brain functions harder, or even impossible to control. This can be due to individual brain damage, but could also stem from a normal variation among people. Some brain functions might also be more intuitive for certain BCI applications. Accordingly, an effort should be made to evaluate new BCI control paradigms in order to adapt to as many physiological and mental limitations, as well as BCI applications, as possible.

Since the use of ECoG and of intra-cortical electrodes requires surgery, the possibilities for doing experiments on humans are quite limited. The main clinical use of ECoG is seizure and functional mapping during presurgical evaluation of patients with intractable epilepsy. In order to perform ECoG BCI experiments one is therefore dependent on access to these patients during the typically one to two weeks they are monitored, but also on their willingness to participate. Since placement of the grids is determined by the clinical demands, these demands also determine the brain functions available for BCI experiments. There is a large body of research on functional brain topography in neuroscience thanks to fMRI (see *Techniques and Implementations*). For about 20 years fMRI has been used as a non-invasive method to investigate how and where the brain processes different types of information. One brain function that has been well studied using fMRI, and that could be of interest for invasive BCI, is the top-down (attentional) regulation of the visual system. One of the most effective and versatile BCI applications would be controlling a computer cursor since it would directly give access to many available software products. Visual attention would be

a very natural candidate for this type of navigation, but since EEG does not have the capacity to distinguish the involved brain regions it has never been investigated.

1.4 Covert Visuospatial Attention

The visual system is located in the occipital lobes, the most posterior part of the brain. The visual cortex is divided into several areas based on location, microscopic structure and function [47]. Different areas process different properties of the visual input, for instance color, motion, orientation and spatial location. The spatial properties of an image are very important, and there are several visual areas processing this information. These areas are ordered in a hierarchy, according to their connections and how early in the information processing chain they are acting. The visual area that first receives the information is called V1, or the primary visual cortex, and is located in and around the calcarine sulcus. Surrounding V1 are the areas referred to as V2 and V3 (Figure 1.2). The retinotopic mapping that transforms the retinal input to the visuocortical areas preserves the spatial organization of the image. V1-V3 are located in both hemispheres; the left hemisphere encodes the visual information coming from the right hemifield and vice versa. Further, the upper and lower hemifields are processed below and above the calcarine sulcus, respectively. The distance from the fovea is represented in the V1-V3 field maps by the distance from the occipital pole. Other, much smaller, visual areas that are involved in the processing of spatial information, e.g. V3A and V3B, are located in the cortex from the anterior part of V3 towards the parietal lobes.

In order to cope with the huge amount of information flowing in via our sensory systems, we have developed the ability to bias the brain and give priority to the sensory features we are interested in. This is referred to as selective attention. When you are scanning the desk to locate your keys, your brain is biased towards the visual features of your keys, such as color and shape, making you more likely to react when they enter your visual field. When you attend to a specific region of your visual field, you will react faster and more efficiently to any stimuli appearing there. In fact, by covertly (without shifting the gaze) attending to a region you make the neurons mapped to it more sensitive and increase their firing rate. It has been well proven [48–55] that the effect of covert visuospatial attention matches the retinotopy and induces activity mainly in the same parts of the cortex as a visual stimulus at the attended region.

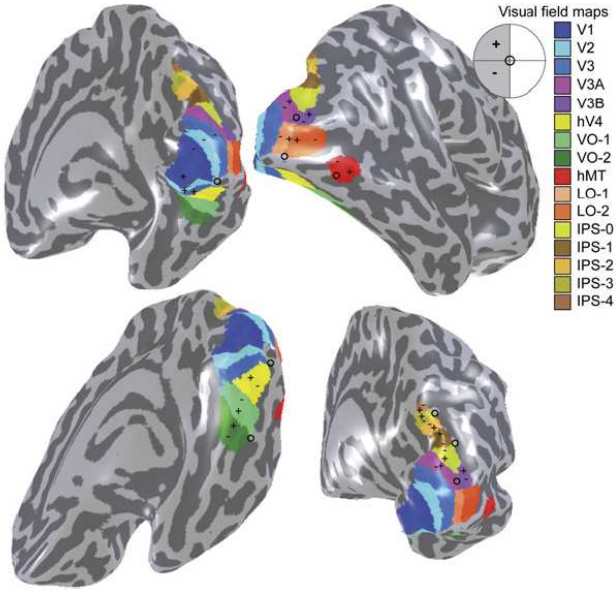

Figure 1.2: The positions of sixteen field maps displayed on an inflated rendering of the cortical surface of a right hemisphere. Fovea and upper/lower visual fields are indicated by the 'o', '+', and '-' symbols, respectively. (Adapted from Wandell et al. (2007) [47])

1.5 Techniques and Implementations

1.5.1 BOLD fMRI

"My hand moves because certain forces - electric, magnetic, or whatever 'nerve-force' may prove to be - are impressed on it by my brain. This nerve-force, stored in the brain, would probably be traceable, if Science were complete, to chemical forces supplied to the brain by the blood, and ultimately derived from the food I eat and the air I breathe."

Lewis Carroll - Sylvie and Bruno, 1889 [56]

Although the quote above precedes fMRI by a hundred years [57–60] it is a surprisingly accurate description of the technique. While the signals in the electrophysiological methods are based directly on the ionic currents present in signalling neurons, functional magnetic resonance imaging (fMRI) instead measures some aspect

of the metabolic processes associated with the neural activity [61]. The most common form of fMRI is based on the so called Blood Oxygen Level Dependent (BOLD) contrast [62].

Active neurons consume a lot of energy in the form of glucose and oxygen. To feed the increased metabolism, more oxygenated blood is sent to the active brain region. In fact, the area is overcompensated creating a net increase of oxygenated blood. The oxygen is transported by hemoglobin and it is the magnetic properties of the hemoglobin that is the basis of the BOLD contrast. Hemoglobin contains iron that gives it paramagnetic properties. However, the magnetic susceptibility of deoxygenated hemoglobin is higher than when it is oxygenated [63]. By using an MRI sequence with ${T_2}^*$ contrast (sensitive to local microscopic inhomogeneities in the magnetic field) images can be produced where the intensity increases with the blood oxygenation level [64, 65]. This is called the BOLD effect while the change in the MRI signal is referred to as the hemodynamic response (HDR). The complete chain of events from neural activity leading to the HDR is a complex combination of changes in metabolism, blood flow and blood volume, and still not fully understood. When analyzing fMRI data a simple model of the HDR to a single event of neural stimulation is used [66]. The signal starts to rise about two seconds after the stimulus and it peaks after about five seconds. If the stimulus continues the signal would eventually reach a plateau and stay there until the stimulus stops. After this the signal drops to below the baseline level until it finally returns completely. This period of decreased signal is called the undershoot.

Compared with electrophysiological signals the BOLD signal is a very slow measure. This means that the temporal resolution of fMRI is very low compared with for instance EEG. On the other hand, fMRI can non-invasively provide very detailed information with a spatial resolution in the order of millimeters and access all parts of the cortex, something no other modality can achieve.

Using a high magnetic field strength has several advantages [67–69]. Firstly, the signal-to-noise ratio (SNR) and the activation contrast in the BOLD signal increases with the field strength. Secondly, at 7T the intra-vascular signal contribution from draining veins is virtually eliminated resulting in a higher spatial specificity. However, there are also negative effects from a stronger field, for instance larger susceptibility artifacts and signal dropouts.

In a typical fMRI experiment blocks of time where the subject is stimulated or performs a task are alternated with blocks of rest, while images are acquired every 1-3 seconds. The signal expected in the processing parts of the brain can be estimated

using the HDR model and compared with the signal recorded from all image voxels. After different preprocessing steps whose purpose are to improve the SNR in the data, a statistical analysis is performed to search for task or stimuli activated brain regions [70].

1.5.2 Real-time fMRI

In real-time fMRI (RT-fMRI) the data is analyzed online during the acquisition. In order to keep up with the scanning, the processing of each image volume cannot take longer than the repetition time (TR). Since the TR is usually around 1-3 seconds, this inflicts high demands on all parts of the processing. The idea of real-time fMRI was first described by Cox et al. [71]. Since then many studies have shown that people can learn to self regulate local BOLD activity with the guidance from real-time feedback. Allowing non-invasive high resolution imaging from all parts of the brain, RT-fMRI can offer feedback from very specific regions that can not be achieved with any other non-invasive modality.

RT-fMRI can be applied to BCI but also in applications such as neurofeedback [72] or to achieve a quick presurgical mapping of important functional areas [73]. In neurofeedback the goal of the self-regulation is to induce behavioral effects and treat symptoms like chronic pain [74] or tinnitus [75]. Below, we will discuss some important aspects of data acquisition, preprocessing and classification that should be considered when performing RT-fMRI BCI experiments.

Data Acquisition and Preprocessing

Naturally, RT-fMRI depends on having real time access to the image data. As the interest in the technique has been growing, it is now more common for MRI scanners to have built-in support for this. Theoretically, all the MRI sequences used for regular fMRI can also be applied to real-time experiments. However, in practice there are some issues that should be considered.

Firstly, since the data most likely needs to be moved from the scanner hardware to an external computer, depending on the available data transfer speed it might be desirable to limit the size of the data volume. For the same reason it might be advisable to avoid using a 3D scan sequence such as PRESTO [76]. In a 3D sequence the image volume is reconstructed only after all its data is available. By using a 2D EPI the data can be moved slice by slice as they are reconstructed over the complete TR.

Secondly, it is important that the SNR is as high as possible. In a regular fMRI

experiment, where the complete data set is first acquired and then analyzed offline, a lower SNR can be compensated for by collecting more data samples. In RT-fMRI, where single trials are classified, this is no longer an option and anything that can improve signal quality, for instance, using a high field scanner, is valuable.

Motion Correction

One major source of artifacts in fMRI data is subject motion, and access to online motion correction can therefore greatly improve the result. Even a small head movement can lead to the extraction of brain activity from outside the areas of interest. The first image volume should be kept in memory and used as a template for a rigid registration applied to all the following images. If the brain regions on which the feedback will be based, have been mapped using an earlier fMRI run (see *Voxel Selection*), the new data should be coregistered with the first image volume of the old run.

Detrending

Another issue is the baseline drift present in fMRI signals [77, 78]. The sources of these low frequency variations can be scanner instabilities but they can also be related to physiological processes such as breathing or spontaneous neuronal activity. This signal drift usually varies in strength and direction over the scanned volume, or even from voxel to voxel. Considering this, the signal drift should be estimated and corrected individually for each voxel and not done globally based on e.g. the average image intensity.

In a standard offline fMRI analysis, based on a general linear model (GLM), the most common way to deal with the problem is to include low frequency regressors and remove their associated variance.

In the case of RT-fMRI the correction can only be based on the data available at the time. If the feedback part is done in a separate fMRI run (see *Voxel Selection*) there will be no data to estimate a drift from during the first part of the scan. By keeping the localization and feedback in a single run, the detrending has the data necessary from the start.

Normalization

The intensities of the fMRI data are commonly normalized to zero mean and unit variance (z-scored) to account for variations in signal amplitude. This step is particularly important when combining data from separate fMRI runs. The baseline and standard

deviation can best be estimated using data collected during rest periods. An alternative to the z-score method is to use the percentage change relative the estimated baseline.

Classification

The majority of studies reporting on real-time classification of fMRI data have used region of interest (ROI) based methods. Each defined ROI represents voxels activated during one of the classes. When a new image volume is available and has been motion corrected, the BOLD values inside the ROIs are extracted, preprocessed and then averaged to give a single value per ROI. The ROI associated with the active class will then stand out as having the highest average.

Offline multivariate analysis of fMRI data has shown that techniques such as Fisher's linear discriminant and support vector machines (SVM) can provide a more sensitive decoding of mental states by combining the temporal and spatial information in voxels from multiple locations in the brain [79]. A comparative study of a selection of different multivariate classifiers applied to fMRI can be found in [80]. Only a few RT-fMRI studies have applied multivariate pattern analysis techniques for real-time classification [81–83] (see review in [84]).

Voxel Selection

It is rare that a particular brain function can be located based only on anatomical landmarks. In order to define the ROIs it is therefore almost always necessary to perform a functional mapping. This mapping can be done in a separate localizing fMRI run or in the same run as, but preceding, the classification. If there is a separate localizing run, the data can be processed using standard offline procedures such as a GLM analysis. So far, the studies where the real-time classification was performed with a multivariate classifier have employed a univariate voxel selection method. A better selection of voxels could most likely be achieved with a multivariate selection method that uses the same algorithm as the classifier [85].

In order to go directly from the localization to classification, without any time in between, the localization must be performed online during the acquisition. One way to do this is by applying incremental algorithms [71, 86–88] that update the results when new data becomes available.

1.6 Outline of this Thesis

In this thesis we propose and evaluate a BCI control task based on visuospatial attention that is completely independent of any stimulation or eye movements. A long-term goal is to implant ECoG electrodes in paralyzed patients for BCI, and the proposed control task should therefore be evaluated for this modality. However, since it is rarely relevant in the presurgical assessment of epilepsy patients to place ECoG grids on the visual cortex, the assessment required an alternative method. Accordingly, we implemented a BCI based on ultra-high field fMRI and used it to test the control task on healthy volunteers. The BOLD changes measured with fMRI have been shown to correlate well spatially with the signals from ECoG [43, 89–91], especially in the higher frequencies of the gamma band (65-95 Hz). Hence, our results are directly relevant for ECoG BCIs and LIS patients.

In **Chapter 2** we show that it is possible to decode single trials of visuospatial attention to two directions, left and right. The left and right peripheral attention target areas contained checkered patterns scrolling in opposite directions. During the first part of the experiment we locate the relevant brain regions using an incremental statistical analysis. During the second part we apply a univariate classification based on the average signal inside the located regions and give real time performance feedback to the subject. We also report on decodability of ECoG signals recorded from part of the visual cortex in a single patient undergoing neurosurgery for epilepsy.

In **Chapter 3**, based on the results from Chapter 2, we add two more directions and decode single trials of left, right, up and down attention. We also remove the scrolling patterns in the attention targets. The subjects receive real time feedback of the performance as in Chapter 2.

In **Chapter 4** we describe an application of visuospatial BCI that takes advantage of its implicit spatial properties and that would be directly relevant for LIS patients. We let subjects navigate a remotely located robot using visuospatial attention. By shifting the attention left, right or up the robot turns left, right or moves forward, respectively. Live video images from a camera attached to the robot make it possible for the subject in the MRI scanner to navigate around a track containing targets that should be reached in a particular order. In contrast to the two previous chapters we here apply a multivariate classification. A linear Support Vector Machine is trained on data acquired during instructed attention and then used to decode the directed attention during a robot control phase.

Chapter 5 attempts to bring the results from previous chapters closer to the ECoG modality. The data from Chapter 3 is reanalyzed offline with different restrictions on

which parts of the cortex are allowed to be used in the decoding. All the restrictions are motivated by issues that need to be taken into account when implanting surface electrodes. We here apply the same type of Support Vector Machine as in Chapter 4.
Chapter 6 gives a summary of the thesis and proposes ideas for future work.

Real-Time Decoding of Visuospatial Attention to Two Directions

Based on: P. Andersson, J.P.W. Pluim, J.C.W. Siero, S. Klein, M.A. Viergever and N.F. Ramsey, "Real-Time Decoding of Brain Responses to Visuospatial Attention Using 7T fMRI", PLoS ONE, 6(11):e27638 (2011)

Abstract

Brain-Computer interface technologies mean to create new communication channels between our mind and our environment, independent of the motor system, by detecting and classifying self regulation of local brain activity. BCIs can provide patients with severe paralysis a means to communicate and to live more independent lives. There has been a growing interest in using invasive recordings for BCI to improve the signal quality. This also potentially gives access to new control strategies previously inaccessible by non-invasive methods. However, before surgery, the best implantation site needs to be determined. The blood-oxygen-level dependent signal changes measured with fMRI have been shown to agree well spatially with those found with invasive electrodes, and are the best option for pre-surgical localization. We show, using real-time fMRI at 7T, that eye movement-independent visuospatial attention can be used as a reliable control strategy for BCIs. At this field strength even subtle signal changes can be detected in single trials thanks to the high contrast-to-noise ratio. A group of healthy subjects were instructed to move their attention between three (two peripheral and one central) spatial target regions while keeping their gaze fixated at the center. The activated regions were first located and thereafter the subjects were given real-time feedback based on the activity in these regions. All subjects managed to regulate local brain areas without training, which suggests that visuospatial attention is a promising new target for intracranial BCI. ECoG data recorded from one epilepsy patient showed that local changes in gamma-power can be used to separate the three classes.

2.1 Introduction

In any interactions with our environment, including speech, we fully depend on the motor system. Damage to neurons involved in motor control can restrict this ability or even completely disrupt communication between our mind and our environment, as in the case of locked-in-syndrome [2]. Situations such as loss of motor function in severe paralysis would greatly benefit from additional means of interaction. By measuring cortical activation changes and linking these changes to commands one can "outsource" the muscular control to a computer and create new channels through which intentions can be transmitted. These techniques are commonly referred to as Brain-Computer-Interfaces (BCI) [92, 93].

Because of its availability and non-invasiveness EEG has been the predominant modality in BCI research. To reach the extra-cranial electrodes the neural electrical potentials have to go through the cerebrospinal fluid, dura mater, skull and scalp. In effect, the signals lose power, bandwidth and spatial resolution. By implanting electrocorticographic (ECoG) or intracortical microelectrode arrays one can record signals much more specific in both time and space, and with a much higher signal-to-noise ratio (SNR), compared to EEG. Encouraged by the success in non-human primates [36, 94–96], there is a growing interest in applying intracranial technologies for human BCI [38, 40, 43, 97–99]. Because the dominating modality in BCI research has been EEG, the control strategies investigated, also for invasive measurements, have mainly been based on systems located in cortical areas accessible by scalp electrodes. The most common strategies have been P300 responses [22–24], steady state visual evoked potentials (SSVEP) [16, 18, 20] and motor imagery [38, 40, 100]. While these types of control have been shown to work in both healthy subjects and patients, many studies have reported that part of the study population is not able to learn control even after training [28, 29, 101–103]. Moreover, patients might have clinical issues making these strategies inapplicable. It is for example uncertain whether paralysed people are capable of engaging their motor cortex after a long period of non-use [30]. This indicates that in the light of intracranial solutions, alternative avenues, using other brain systems, are worth exploring to further the BCI field and to be able to create an individually optimized setup for each patient. While the term "BCI illiteracy" is sometimes used for subjects not able to control a BCI it is more likely that the particular control task is not suitable and that by choosing the right task also these subjects can learn to gain control.

Intracranial electrodes make it possible to access brain functions that are located deeper in the brain or are otherwise inaccessible for EEG.

Here we present a new avenue for intracranial BCI, which exploits specific properties of the visual system. With the help of attention we can select what sensory information to focus processing resources on [104]. Covert visuospatial attention, i.e. focusing attention on a specific part of the visual field in order to better process what happens in this spatial region while maintaining gaze at the center of the field, is known to induce changes in activity in the visuospatial cortex [48–52, 105].

Two earlier studies have examined the potential use of brain activity associated with covert visual spatial attention for BCI control, and demonstrated that changes in the alpha band could be detected using MEG [106] or EEG [107]. The induced changes could be classified with offline techniques, but as realtime analysis was not tested it is not clear whether these changes offer enough detail for BCI application. Spatial detail of MEG and EEG may be a limiting factor in exploiting the brain activity patterns associated with covert attention. The retinotopy assures that attention to a restricted part of the visual field corresponds to activity in restricted cortical areas. An intracranial BCI system with high resolution, based on e.g. ECoG, should be able to react only to local attention and not to attention anywhere in the peripheral visual field. More attention target regions can be added to increase the degrees of freedom. Moreover, an attention target region could be moved to the place in the visual field mapped to the cortical region most suitable for implantation.

In principle, functional MRI yields better detail compared to MEG and EEG, at the expense of speed of detecting changes. Moreover, it is inherently sensitive to activity anywhere in the brain, and as such can be used to investigate new alternative control tasks and cortical regions. Real-time fMRI [8, 9, 74, 108, 109] offers the possibility of identifying target regions for intracranial electrode placement presurgically and can be used to train the patient beforehand. Although fMRI measures bloodflow as opposed to electrical or magnetic signals, fMRI activations have been shown to agree with those found with ECoG [43, 89–91]. Spatial correlations between activity patterns obtained with both has been shown to be particulary strong in the high gamma band (> 60Hz) [91]. The use of real-time fMRI for learned self-regulation of local brain activity has been demonstrated several times before [9]. Most of these studies have had a neurofeedback approach, where the self-regulation was not investigated with the purpose of transmitting commands. Here the feedback was given directly on changes in the BOLD signal. Building on these results, the technique has also been applied for BCI purposes where the signal changes are classified to discrete outputs representing intentions (see review in [110]). Activity induced by covert attention is rather subtle

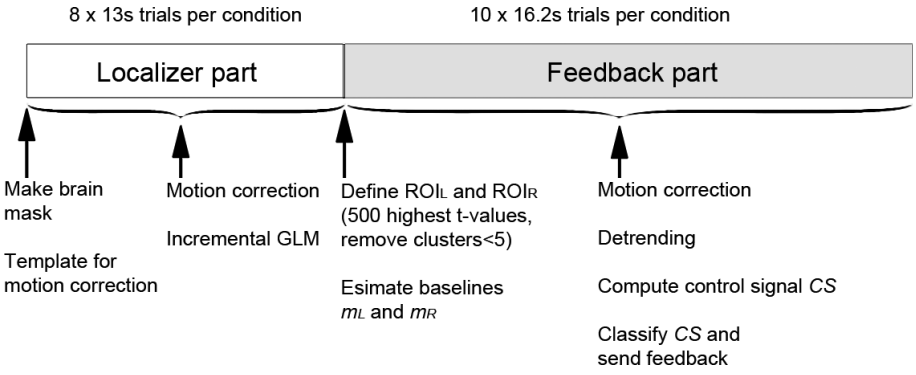

Figure 2.1: Illustration of the experiment timeline. The localizer and feedback data are acquired in the same fMRI run.

and for the present purpose of real-time decoding, requires the most sensitive fMRI technique available. Ultra-high field MRI systems have become available recently, and have been shown to yield excellent sensitivity [67]. To test our hypothesis we implemented real-time fMRI on a 7 Tesla MR scanner using healthy volunteers. We postulate that if real-time decoding is feasible with covert attention and fMRI, placement of electrodes on the visual cortex should also yield decodable signals. We also report on decodability of ECoG signals obtained from visual cortex in a patient undergoing neurosurgery for epilepsy.

2.2 Materials and Methods

The experiment was performed in a single fMRI run in which the healthy volunteers were instructed via a central cue to move their attention to one of three target regions while maintaining their gaze at the center. The scan consisted of two parts; a first part in which we located the activated regions and a second part in which subjects were given real-time feedback based on the activity in these regions. An overview of the full experiment can be seen in Figure 2.1.

2.2.1 Subjects

fMRI data were acquired from ten healthy volunteers (age 19-27, 6 female, all except one right handed). One of the subjects showed very poor performance during the experiment. After the experiment the subject communicated problems with concentration and offline inspection of the fMRI data showed excessive motion. Based on this we have excluded this subject. Two additional subjects performed the task outside the scanner while we recorded their eye movements using electrooculography (EOG). Multi-channel subdural ECoG data was recorded from one patient (female, age 26, left hemisphere) undergoing neurosurgery for epilepsy.

The protocol was approved by the ethics committee of the University Medical Center Utrecht in accordance with the declaration of Helsinki (2008), and all subjects had given their written informed consent. All subjects were naive to the task.

2.2.2 fMRI data acquisition and real-time system

The data were collected on a 7T Philips Achieva system with a 16-channel headcoil. The functional data were recorded using an EPI sequence (TR/TE=1620/25ms; FA=90; SENSE factor=2; 35 coronal slices, acquisition matrix 96x96, slice thickness 2mm with no gap, 1.848 mm in-slice resolution). The FOV was selected so it covered the occipital lobe. A total of 500 volumes were acquired in a single run and divided into 200 volumes of localizing relevant brain areas (localizing part) and 300 volumes of real-time feedback based on activation in these located regions (feedback part). Directly following reconstruction on the scanner the data were sent to a separate computer performing the analysis (Dual-Core 2.5GHz notebook) via the local network using a TCP/IP protocol and the Philips DRIN (Direct Reconstruction INterface) module.The stimulus was projected to the subject from a second computer via a video projector. An update-trigger containing information about the direction and color of the instruction marker was sent to the second computer via a serial cable. Except for the motion correction all the parts were implemented in Matlab (Mathworks, Natick, MA).

2.2.3 Task

The visual stimuli were constructed as two rectangular areas, one in the left peripheral visual field and one in the right, each containing a checkered pattern and both at a vi-

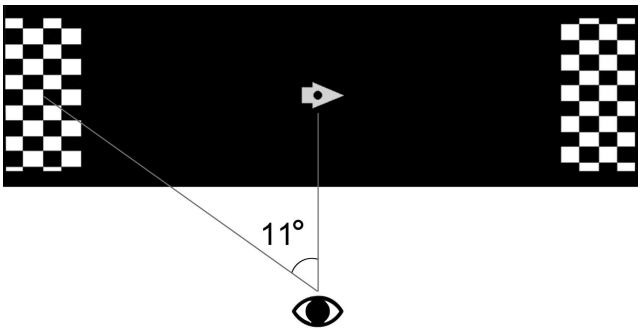

Figure 2.2: The visual stimuli.

sual angle of 11 degrees relative to a central cue (Figure 2.2). To facilitate the shifting of attention direction, we made the checkered patterns scroll (2s per cycle) upwards on the right side and downwards on the left. In the center was a marker on which the subjects were instructed to fixate their gaze at all times. Both checkerboards were constantly visible throughout the fMRI runs, while the center marker was alternated between a right arrow, a left arrow and a circle. The arrows indicated to which side the subject had to direct the visual attention. The circle indicated that the attention was to be directed to the center. The three trial types were repeated in a pseudo random scheme with the restriction of no two adjacent attention trials being in the same direction.

2.2.4 Localizing part

Trials

The localizing part consisted of eight trials of each condition plus one extra initial central attention trial, each being eight scans (13.0 seconds) long. The instruction was updated first after the analysis finished (1.0 seconds on average). This time has been accounted for in all plots and results.

Motion correction

The first volume was used as the template for motion correction and all the subsequent volumes were aligned to it using a rigid transformation. The registration was performed by minimization of the sum of squared differences between grey-value

intensities. To achieve real-time performance, a stochastic gradient descent method [111] was employed for optimization, using 50 iterations. The images were blurred with a Gaussian filter (σ = 1 voxel) prior to image registration. Linear interpolation was used during optimization while cubic B-spline interpolation was used to generate the final rotated/translated image. The algorithm was implemented in C++, and called from Matlab. The computation time was approximately 0.6s per fMRI volume.

Analysis

To find the activated voxels in real time we implemented the incremental GLM method described in [87]. The incremental approach ensures that the computation time does not grow with the number of scans. By keeping the whole experiment in a single run we minimize the risk of movement between selection of ROIs and the feedback experiment and we get an improved estimation of the low frequency drift and therefore a better detrending and a better control signal. Three regressors representing right- and left-sided attention, and a linear function as a simple model for the drift were included in the model. Since visual spatial attention induces both increased BOLD signal in retinotopically mapped regions and decreased signal in unattended regions [50, 55, 112, 113], the differential contrasts "right-left" and "left-right" were used when computing the t-maps. This also made sure we avoided picking up regions responding to attention in general.

ROI selection

When the localization part was finished (200 volumes) the resulting t-maps were used for making the two sets of voxels representing right versus left side attention and left versus right side attention (denoted $\mathrm{ROI_R}$ and $\mathrm{ROI_L}$ respectively) as follows. First the two t-maps were masked to only include values inside the brain. The mask was constructed by first thresholding a smoothed image volume and then filling any holes. The two most anterior of the coronal slices were excluded from the mask to exclude boundary artifacts from the registration. For each of the two t-maps the voxels with the 500 highest t-values were selected and from these clusters smaller than five voxels were removed. The remaining sets of voxels constituted the $\mathrm{ROI_R}$ and $\mathrm{ROI_L}$. Next, a baseline value was computed for each ROI, m_R and m_L, by averaging the signal inside the ROI in the data recorded during the central condition. The first three volumes (4.86s) in central trials that were preceded by an attention trial were excluded to let the signal return to baseline. Additionally, the individual time series of the voxels

making up the ROIs were saved for the purpose of detrending during feedback.

2.2.5 Feedback part

Trials

During feedback a longer trial of 10 scans (16.2 seconds) was used, and each condition was repeated 10 times. Feedback was given by coloring the central instruction marker according to the performance (see *Classification and feedback*). As during the localizer part, the instruction was updated after the analysis (0.8 seconds on average). Also here we have accounted for this delay in all results.

Analysis

In the feedback part of the scan we gave the subjects real-time information about their performance based on the activity in $\mathrm{ROI_R}$ and $\mathrm{ROI_L}$, as follows. When a new volume was available it was first motion-corrected as during the localizing part. After this the values inside the two ROIs were extracted and added to the time series of available data (including the localizer part). To remove any low frequency drift [77] in the signal, detrending was now applied using an algorithm originally described in the context of real-time detrending of heart-rate variability measurements [114] ($\lambda = 200$). Each voxel's time series was detrended individually since the signal drift looked quite different in different parts of the image. The new detrended values were averaged to give a single value per ROI and fMRI volume (k), $s_R(k)$ and $s_L(k)$. These numbers were in turn normalized to a percentage change from the baseline and subtracted to give the value of the control signal CS defined as

$$CS(k) = 100 \times [(s_R(k) - m_R)/m_R - (s_L(k) - m_L)/m_L] \tag{2.1}$$

where k is the volume number and m_L and m_R are the baseline values computed from the localizer data.

Classification and feedback

The control signal was classified based on its magnitude using three thresholds above the baseline ($\mathrm{p_1}$, $\mathrm{p_2}$ and $\mathrm{p_3}$) and three below ($\mathrm{n_1}$, $\mathrm{n_2}$ and $\mathrm{n_3}$). The central instruction marker was then colored according to this classification. During attention two tones of

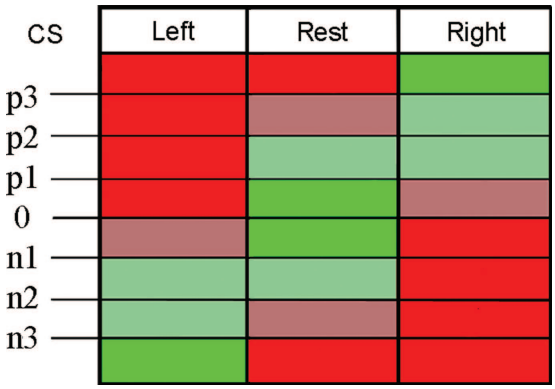

Figure 2.3: Table displaying the colors used for performance feedback.

green represented weak and strong signals in the correct direction whereas two tones of red represented a control signal indicating the wrong (or lack of) direction (see Figure 2.3). During the central condition green represented a signal close to baseline.

For Subjects 1-7 fixed CS thresholds of p_1=1.5, p_2=2.5, p_3=4, n_1=-1.5, n_2=-2.5, n_3=-4 were used. These values turned out to be rather conservative, and for Subjects 8-9 an adaptive thresholding approach was applied, where the localizing data were used to select individual values online. First a retrospective CS was computed applying Equation 2.1 to the available (localizer) data. Then, for both right and left attention, the thresholds required to limit the false positive rate (FPR) to 0.2 were estimated. These estimated thresholds were used as p_1 and n_1. Here we needed a binary classification and for right attention the value of CS was classified as 'positive' if larger, and 'negative' if smaller than p_1. Thus, the FPRs were computed using false positives from both the other conditions, i.e. opposite and center attention. In order to account for the hemodynamic delay, the instructions were shifted 3 TRs with respect to the control signal before computing the FPR. The other threshold levels were now set as $p_2 = 3 \cdot p_1$, $p_3 = 4 \cdot p_1$, $n_2 = 3 \cdot n_1$, $n_3 = 4 \cdot n_1$.

2.2.6 Performance

The true positive rate (TPR) and false positive rate (FPR) were used as a measure of performance. As when determining the adaptive thresholds, the instructions were

shifted 3 TRs to account for the hemodynamic delay. The FPR was computed both including and excluding the central condition. The reason for considering only the attention blocks is that the BOLD undershoot following an attention block may produce a rebound in CS towards the opposite side of the baseline. This is a BOLD effect and would not be present in a BCI based on electrophysiological measurements, e.g. EEG. Hence, to give a fairer measure of stability during attention, the FPR was also computed after removing the 'attend center' blocks. To visualize how the TPR and FPR depended on the thresholds, they were computed for varying threshold levels and the results were plotted as receiver operating characteristic (ROC) curves.

The performance depends on the thresholds p_1 and n_1, and since only Subjects 8 and 9 were classified using adaptive thresholding, we also recomputed the performance for Subjects 1 to 7 offline applying the same adaptive method.

2.2.7 Offline group analysis

For the group analysis we used SPM5. Each subject's realigned data were normalized to the Montreal Neurological Institute (MNI) space using the structural T1 image. The normalized functional images were smoothed with an isotropic 4mm FWHM Gaussian kernel and then used to compute activation patterns. The second level analysis was performed using a paired t-test (attend left, attend right) on the resulting beta images and the contrasts right-left and left-right were applied.

2.2.8 ECoG data and analysis

The patient had a 64-channel (8x8) electrode grid positioned on the left parietal-occipital cortex, covering a considerable part of the cortex included in the fMRI volume for the healthy volunteers. Data were collected during a localizer task (20 trials attend left, 20 attend right, 39 attend center, no feedback), with 5 s trial duration. The signal was acquired at 512 Hz, and was referenced to a common average across all 64-channels. The first 4 seconds (after instruction) of each trial were used to compute the power in the high gamma band (65-95 Hz). This single band was chosen as fMRI matched this frequency range in previous studies [43, 91, 115]. Performance was estimated by means of a leave-one-out cross-validation approach. For each trial, all the other trials (constituting a "training set") were used to create a classifier. Each classifier was a simple linear combination of channels (electrodes), resembling the fMRI approach. Each of the 79 cross-validation tests was performed in three steps;

(1) Normalization, (2) Selection of channels and (3) Classification of the test trial.

(1) All channels (including the test trial data) were normalized, to zero mean and unit variance, as estimated using the 78 training trials.

(2) Three sets of channels were identified, one for each attention direction. Each set contained the channels where the average amplitude differed enough between the corresponding direction and the other two. For a channel to be included in one of the sets it should; a : have an average higher (lower, if negative due to deactivation) during this attention direction than for the other two, with a difference to the closest one larger than a certain threshold δ_d (see *Optimization*). b : have an average during this attention direction with a value exceeding half the standard deviation computed over all three directions. Note that a channel can be selected for two of the attention directions if there is an increase in gamma power during one direction, and a decrease in the other.

(3) The test trial was classified as the direction whose set of channels had the highest average magnitude.

Optimization

To optimize the performance, the selection of channels was computed for a range of thresholds ($\delta_d \in [0.1, \quad 0.4]$), each giving a different selection. For each of these selections the training trials were themselves classified, as described above. The final selection, used for classifying the test trial, was the one giving the most correct classifications of the training data.

2.3 Results

2.3.1 Control signals

The control signals (CS) for all subjects are plotted in Figure 2.4. Dark gray, light gray and white represent the left attention, right attention and center conditions, respectively. The condition blocks are shifted 3 TRs (4.9s) to compensate for the hemodynamic delay. The responses to the different conditions were also averaged, first for the individual subject, then over all nine subjects. The results are plotted in Figure 2.5. The strength of the attention-modulated signal changes in ROI_R and ROI_L relative to their baselines might not be equal. This means that when subtracted (see Equation 2.1), CS may be biased towards one of the directions. Such an effect can be seen in

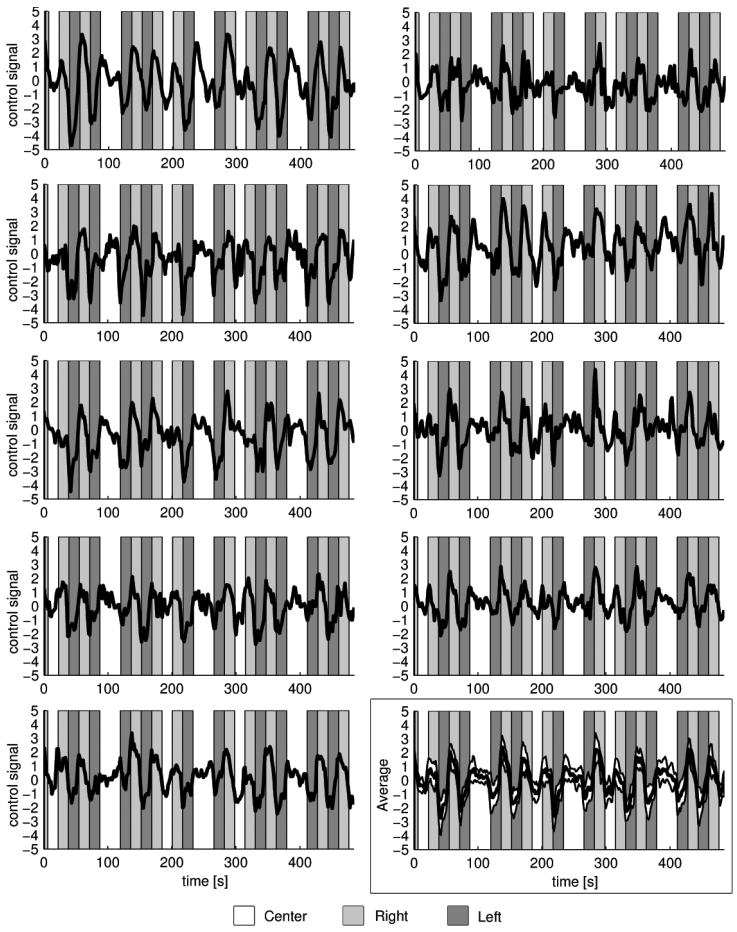

Figure 2.4: The control signals (CS) for all subjects. (Subject 1-9 from left to right and top to bottom.) Light and dark gray represent right-sided and left-sided attention respectively. The blocks have been shifted 3 TRs (4.9s) to compensate for the hemodynamic delay. The last plot shows the average control signal over all subjects, with the standard deviation shown in white.

Figure 2.4 for Subjects 3 and 5. This bias can in turn lead to a difference between the two sides in the time needed to exceed the thresholds. A more laterally symmetric control signal, and one that is more uniform across subjects, could be achieved by normalizing the signals using both the baseline and the standard deviation as $(s_k - m)/\sigma$ (see Equation 2.1).

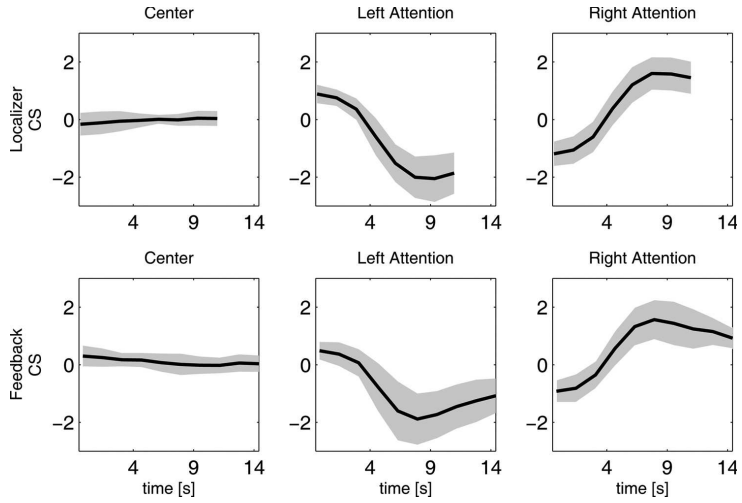

Figure 2.5: Average control signal during the central, right-sided attention and left-sided attention trials. The averages are shown both for the actual control signal during feedback and the control signal computed offline using the localizing data. The standard deviation is shown in gray.

2.3.2 Performance online

Table 2.1 shows the true positive rates (TPR) and false positive rates (FPR) from the online results. For Subjects 1 to 7 fixed CS thresholds of $p_1 = 1.5$ and $n_1 = -1.5$ were used. The low number of true positives together with the near absence of false positives indicates that these thresholds were rather conservative.

For Subjects 8 and 9 adaptive thresholding was applied. The localizer data were used here to estimate what threshold levels are needed to restrict the FPR to 0.2. In this way we could increase the number of true positives, while estimating the risk.

2.3.3 Performance offline

To assess what the performance would have been if we had applied the adaptive thresholding to all subjects we recomputed the analysis offline for Subjects 1 to 7. In this analysis the thresholds were based on the localizer data in the same way as was done online for Subjects 8 and 9. Table 2.2 shows the new thresholds together with the resulting TPR and FPR values.

A more detailed view of how the classification depends on the thresholds is given by

Table 2.1: Online performance. Online True Positive Rate (TPR) and False Positive Rate (FPR). Fixed thresholds of 1.5 and -1.5 were used for Subjects 1 to 7, whereas adaptive thresholding was applied to Subjects 8 and 9 (see Table 2.2). (L = left attention, R = right attention, † excluding the 'attend center' condition.)

	TPR		FPR		FPR†	
Subject	**L**	**R**	**L**	**R**	**L**	**R**
1	0.64	0.39	0.02	0.07	0.00	0.00
2	0.28	0.13	0.03	0.00	0.00	0.00
3	0.55	0.08	0.04	0.01	0.01	0.00
4	0.17	0.54	0.02	0.09	0.00	0.02
5	0.68	0.21	0.00	0.00	0.00	0.00
6	0.18	0.29	0.02	0.01	0.00	0.00
7	0.40	0.11	0.01	0.00	0.00	0.00
Avg	0.41	0.25	0.02	0.03	0.00	0.00
8	0.73	0.81	0.19	0.29	0.09	0.13
9	0.77	0.89	0.18	0.31	0.05	0.16
Avg	0.75	0.85	0.19	0.30	0.07	0.15

Table 2.2: Offline performance. Offline True Positive Rate (TPR) and False Positive Rate (FPR) for Subjects 1 to 7 when applying adaptive thresholding. Columns 2-3 show the corresponding thresholds. We also included these numbers for Subjects 8 and 9 where the method was applied online. (L = left attention, R = right attention, † excluding the 'attend center' condition.)

	threshold		TPR		FPR		FPR†	
Subject	**L**	**R**	**L**	**R**	**L**	**R**	**L**	**R**
1	-0.27	0.22	0.89	0.78	0.28	0.20	0.13	0.03
2	-0.32	0.16	0.90	0.67	0.29	0.14	0.13	0.01
3	-0.52	0.20	0.91	0.69	0.21	0.20	0.13	0.00
4	-0.27	0.37	0.61	0.94	0.12	0.36	0.00	0.18
5	-0.34	0.40	0.99	0.53	0.26	0.13	0.26	0.00
6	0.01	0.31	0.82	0.73	0.30	0.23	0.17	0.06
7	-0.06	0.18	0.89	0.88	0.24	0.23	0.09	0.07
Avg			0.86	0.75	0.24	0.21	0.13	0.05
8	-0.15	0.22						
9	-0.24	0.18						

the ROC curves in Figure 2.6. The TPR and FPR were computed using thresholds between -1 and 5 for right-sided and 1 and -5 for left-sided attention. The unit thresholds are marked in the plots as squares and triangles and the levels for the online thresholds, 1.5 and -1.5, are indicated by stars. Note that these values were not used online for Subjects 8 and 9.

An overview of the classification results over the different trials is presented in Figure 2.7. For each time point (not adjusted for hemodynamic delay) it shows the

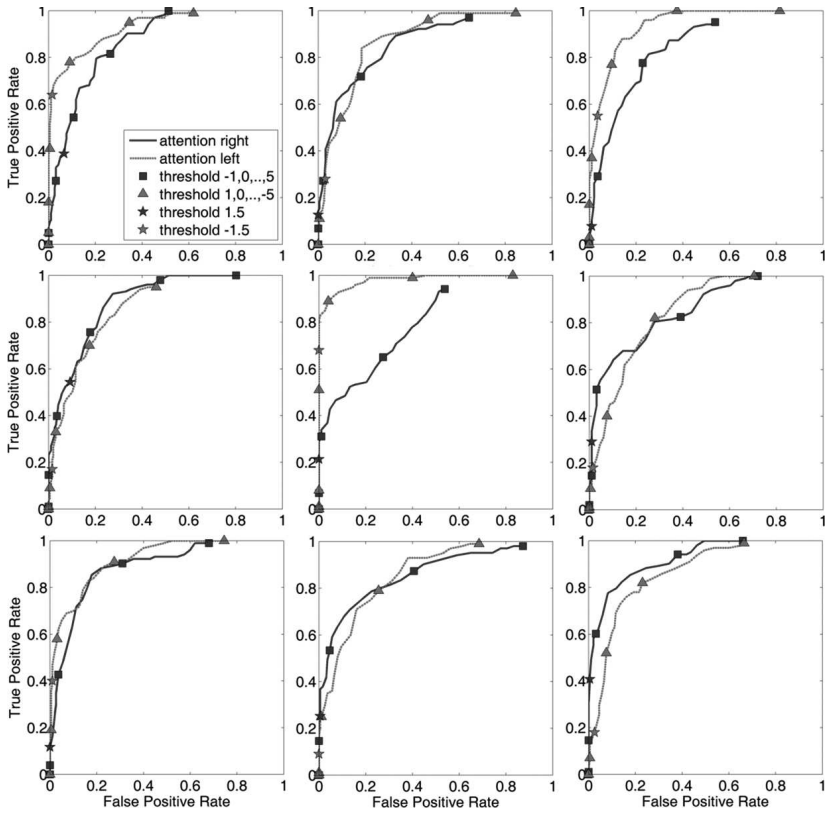

Figure 2.6: ROC curves plotted for the control signal over varying thresholds. (Subject 1-9 from left to right and top to bottom.)

number of subjects with a correct classification. We also computed the percentage of all trials, for all subjects, that would be correctly classified if based on a single volume. The curves in Figure 2.7 show the results for classification based on each of the 10 time points within the trials. Classifying the trials using only the 5th time point gives an average correct classification of 89% for left attention and 88% for right attention.

Though the RT-fMRI setup presented here is not meant to be directly used as a BCI, but rather as a tool to practice and evaluate control tasks, a bit rate can be computed. The most commonly used bit rate definition in the field of BCI is the

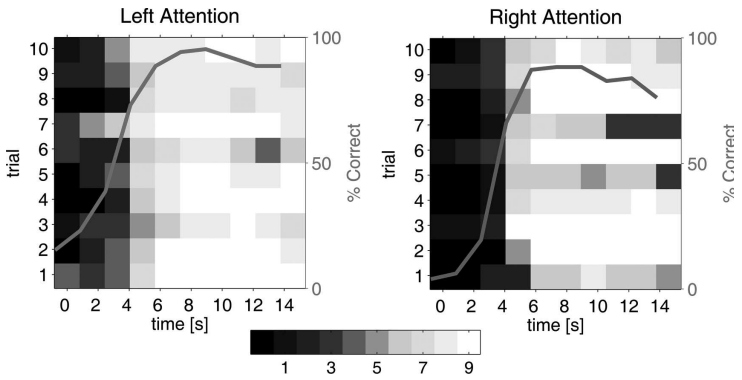

Figure 2.7: Number of subjects having a specific image volume correctly classified. Each row represents one of the 10 trials, and each column a time point (not adjusted for hemodynamic delay) in that trial. The curves show, for all time points, how many trials would be correctly classified if based only on this particular volume.

one from Wolpaw [116]. This definition assumes that the classification accuracy is the same for all classes and that the errors are equally distributed. To fulfil these requirements we excluded the center class so that each left and right attention trial were assigned to either left or right. When each trial was classified using only the fifth time point the average accuracy was 92% (left 93%, right 91%), with the increase due to having no false negatives from the central attention class. With each trial being 16.2 seconds this gave a bit rate of 2.2 bits/minute. This number should not be seen as a highest possible bit rate using a two direction visual attention task. Based on a direct measure, e.g. ECoG, the time needed to make a classification will be much shorter.

2.3.4 ROI selection

The t-maps from the online analysis of the localizing data were thresholded to the two ROIs, ROI_R and ROI_L. Off-line inspection showed that the t-values corresponding to the 500 voxels threshold were between 2.56 and 4.85 (ROI_R; mean $= 3.61, \sigma = 0.72$, ROI_L; mean $= 3.67, \sigma = 0.60$). Table 2.3 shows the individual values for both ROI_R and ROI_L as well as the size of the final ROIs, i.e. after removing clusters smaller than 5 voxels.

Table 2.3: T-value thresholds and ROI sizes. T-values corresponding to the threshold of 500 voxels used to define the ROIs. |ROI| is the number of voxels in the final ROI, after removing all clusters smaller than five voxels.

Subject	t ROI_R	$\|ROI_R\|$	t ROI_L	$\|ROI_L\|$
1	4.71	456	3.89	367
2	2.99	397	3.03	389
3	2.85	293	3.70	389
4	4.49	420	2.85	330
5	4.36	410	4.85	397
6	3.45	439	3.66	425
7	2.56	331	3.03	384
8	3.36	428	4.26	363
9	3.71	297	3.76	260

2.3.5 Incremental GLM analysis

The incremental GLM makes it possible to do the whole experiment in a single fMRI run. The alternative is to stop after the localizer data have been collected to do the statistical analysis and define the ROIs, and then restart to do the feedback part. Offline comparisons show that the incremental method [87] gives an end result very similar to a standard 'full data' GLM analysis using the same regressors and contrasts. ROIs were for the latter method computed as online, but based on t-maps computed from the full localizer data set at once, instead of in incremental steps. These 'full data' ROIs, $\overline{ROI}$, were then compared to the incremental ROIs, *ROI*, using the Dice coefficient computed as

$$DC = \frac{2|ROI \bigcap \overline{ROI}|}{|ROI| + |\overline{ROI}|}$$

where $|\cdot|$ is the volume.
The average numbers across subjects, 0.98 for ROI_R and 0.99 for ROI_L, indicate an almost perfect overlap and suggest that using the online incremental GLM does not decrease the sensitivity.

2.3.6 Group analysis

To find the most frequently activated cortical regions during the two attention conditions a group analysis was conducted. The t-maps from the second-level analysis are displayed in Figure 2.8. The activation patterns for all individual subjects (transformed to MNI space) are displayed in Figure 2.9 both for the localizing and the feedback data. Figure 2.10 shows the group distribution of voxels selected for the

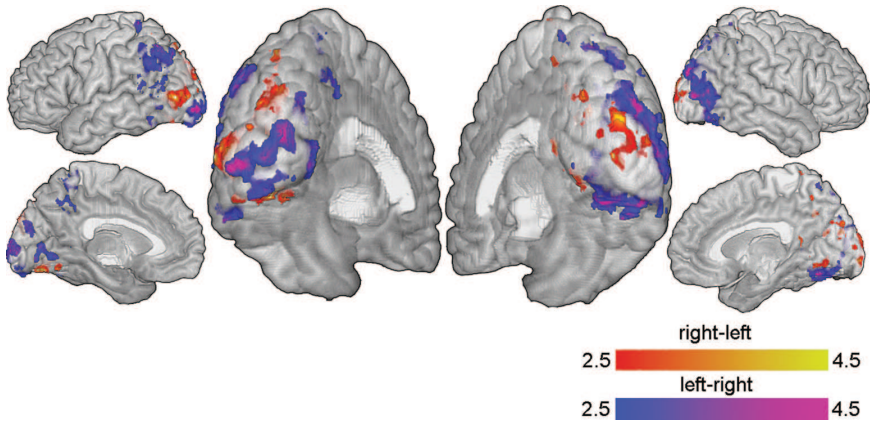

Figure 2.8: The group activation pattern. Red represents t-values from the contrast 'attend right-attend left' while blue represents 'attend left-attend right'.

ROIs, projected on transversal slices.

The contrasts, and therefore the control signal, are sensitive to both activation during attention to one side and deactivation during attention to the opposite side, i.e. a high t-value for 'right-left' can be due to increased activity during right attention or decreased activity during left attention, or both. Figure 2.11 separates the areas in Figure 2.8 into voxels contributing to the differential contrasts by means of positive activation and voxels whose contributions come from a deactivation during opposing attention. An interesting effect can be seen in the foveal regions around the occipital poles in Figure 2.8 and Figure 2.11. These regions show deactivation during contralateral attention. A possible explanation could be that part of the visual field between center and attended periphery is suppressed to reduce interference.

2.3.7 Eye movements

Any eye movements correlated to the instructions could induce activations falsely interpreted as attention related. If these regions end up in the ROIs it would mean that the regulatory control would partly be based on motor activity. Even though it has been shown multiple times that people have no trouble performing covert spatial attention shifts in the absence of any eye-movements (e.g. [49, 51, 106]) we decided to test subjects' abilities to perform the task while maintaining a central fixation. Without an eye-tracker approved for use at 7T, we could not record the eye movements during the

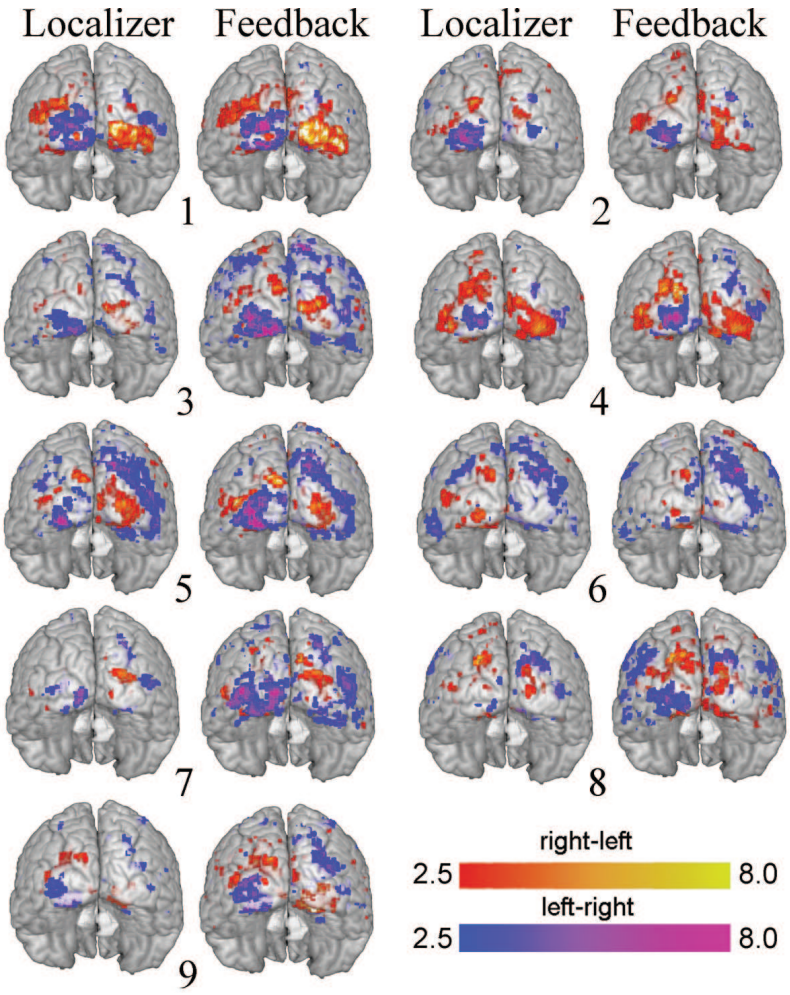

Figure 2.9: The individual subjects' activation patterns. The patterns both during the localizer part and the feedback part are displayed on the MNI brain. The red and blue color scales represent t-values from the contrasts 'attend right minus attend left' and 'attend left minus attend right'.

experiment. Instead we had two additional subjects, naive to the task and not part of the rest of the study, performing the task outside the scanner during which we recorded their eye movements using electrooculography (EOG) with two electrodes below and

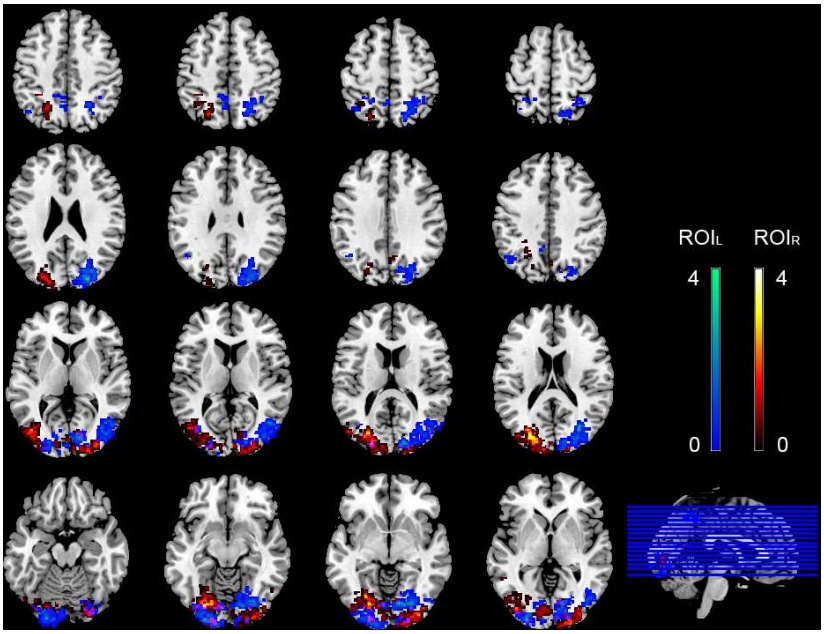

Figure 2.10: The number of subjects having a voxel included in an ROI. The red scale represents $\mathrm{ROI_R}$ and the blue scale $\mathrm{ROI_L}$. Due to interpolation during normalization, the numbers are not integers.

lateral to the right eye, and a reference electrode behind the ear. These subjects showed no eye movements correlated to the task. Figure 2.12 shows the average EOG response in both electrodes for one of the subjects.

The activity patterns themselves can also be used as an indication of whether or not eye-movements were present. If the gaze is moved to fixate on one of the targets, this target will move to the center of the visual field while the instruction cue, to which the subject will now have to move the attention in order to notice new instructions, and the opposing target will be located in the contralateral hemifield. Since each hemifield is represented by the contralateral visual cortex this would mean that, except for the foveal region, only the ipsilateral side would be activated. In other words, if the subjects moved the gaze to the targets instead of keeping it fixed at the center, the "left attention" condition would only show activity in the right hemisphere and vice versa. This was confirmed for Subject 8 in an additional localizing run where the subject was asked to move the fixation the checkerboards. When compared to

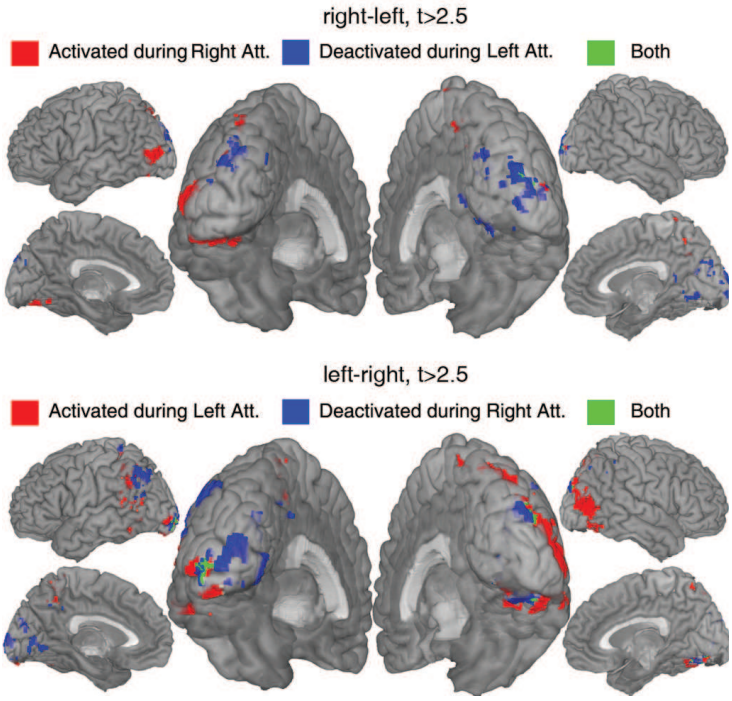

Figure 2.11: Activations and deactivations. The group t-values higher than 2.5 (see Figure 2.8) are separated into areas showing activation versus areas showing deactivation relative to the central attention task. The upper half shows the contrast 'attend right-attend left' and the lower half 'attend left-attend right'. Red represents voxels whose contribution comes from increased activity, blue the voxels showing deactivation during attention to the other side, and green voxels showing both these effects.

the pattern seen during covert attention, the result is distinctly different and laterally mirrored (see Figure 2.13). If the subjects instead made small saccades towards the target and back, the BOLD signal changes would not have been strong enough for us to classify them in single images.

2.3.8 ECoG

The average TPR over the 79 cross-validation tests was 0.70 (right: 0.55, left: 0.60, center: 0.82). It should be noted that almost half of the trials were center attention. Figure 2.14b shows the number of times an electrode was selected to be included in the classifier for one of the leave-one-out tests, and for which class. The yellow markers

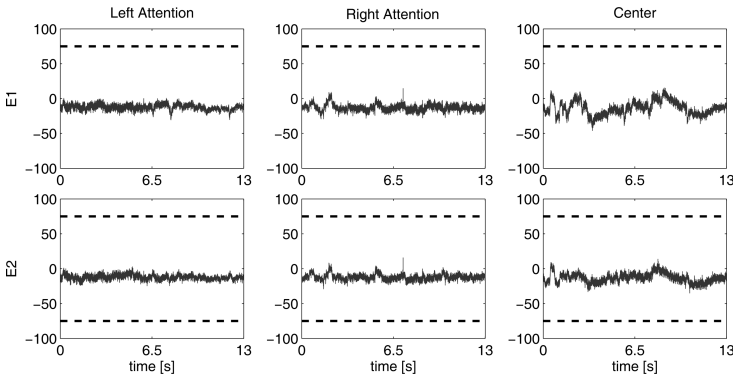

Figure 2.12: The plots show the eye movements for one of the two subjects (not part of the rest of the study) measured using EOG outside the scanner while performing the localizer task. Two electrodes, E1 and E2, were placed below and lateral to the right eye, respectively, and were referenced to an electrode placed behind the ear. The dotted line shows the response level during actual eye movements to the target regions (two lines for E2 since the response to the two directions has opposing polarity.)

show the locations of the electrodes, and the colored circles the selection frequency. The locations of the selected electrodes can be compared to the fMRI groupmap in Figure 2.14a.

2.4 Discussion

In this study we show that brain signals associated with covert visuospatial attention can be used for BCI control. Unique to this approach is that the user can process information in the central visual field while simultaneously exerting control over a device merely by directing attention to the peripheral field. The brain activation patterns confirm earlier studies on visuospatial attention, but are here decoded in real-time. Our subjects easily managed to avoid eye movements during the task. The results, together with the fact that fMRI activations have been shown to agree with those found with intracranial electrodes [43, 89–91], have direct implications for BCI implants. A high performance across subjects and activation confined to a few small brain areas, suggest that the new control paradigm is well suited for intracranial implants.

Using a 7 Tesla MRI system we show that signals from the visual cortex are highly correlated to the direction of visual attention, and can be reliably decoded in real time. One could argue that the use of moving checkerboards introduces a confound

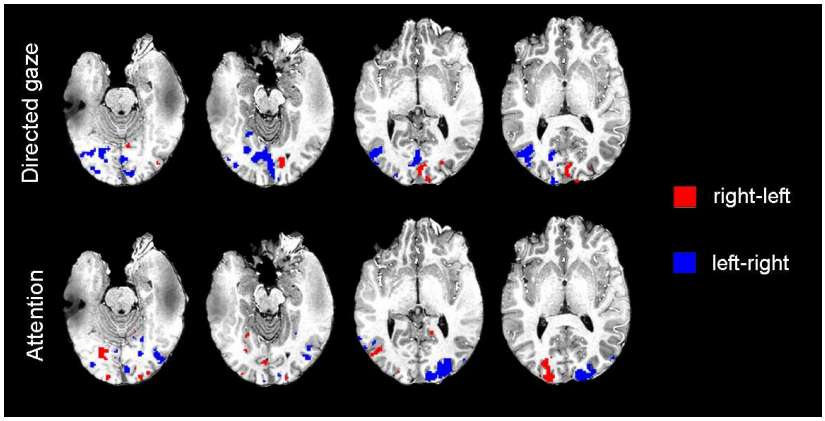

Figure 2.13: The difference in activation pattern between covert attention and actual directed gaze. The localization part of the experiment was repeated for Subject 8 with the instruction to direct the gaze to the target. The overlay show t-values > 3 for the contrasts 'right-left' and 'left-right'.

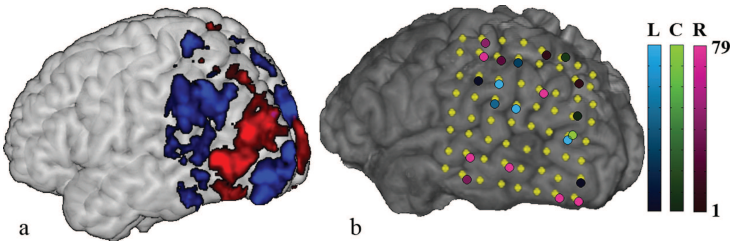

Figure 2.14: ECoG electrode selections. (a) The fMRI group activation pattern ($t > 1.5$, red: 'right-left', blue: 'left-right'). (b) The yellow markers show the electrodes' locations on the cortical surface. On top of the markers it is shown in how many of the leave-one-out tests the electrode was included. Red, blue and green represent right, left and center attention.

by inducing activity due to visual motion (e.g. in area V5). However, we scanned two subjects using stationary stimuli (simple triangles in the periphery) and found the same activation pattern and performance (TPR; 90%/80% and 80%/80% for left/right attention). Further, by using the checkerboard stimuli we show that real visual input would not necessarily affect the attention-based control signal. Thus, even in real-life situations with input covering the full visual field our attention-based BCI approach is likely to work, although this requires further testing.

Cortical activations

Visuospatial attention, i.e. attention to central or peripheral parts of the visual field while maintaining gaze to the center, has been shown to cause region-specific changes in brain activity as measured with fMRI [48–55]. An important finding was the close topographical match of regions activated by actual visual stimuli and by mere attention. Topography of the visual field on the visual cortex has been elucidated in great detail [47]. In V1 and in the encircling areas V2 and V3 each hemifield maps onto the contralateral hemisphere, and stimuli above and below the horizontal meridian are mapped onto the ventral and dorsal regions respectively. From the center of the visual field to the periphery, cortical representations are laid out from the occipital pole towards more anterior aspects of the visual cortex.The parts of the visual cortex that correspond to the attended region exhibit an increased BOLD signal during directed attention, also in the absence of a visual input or eye movements [55, 117]. Importantly, while brain areas processing the attended location exhibit an increased BOLD signal, a decreased signal is seen in brain areas responsible for the part of the visual field surrounding the attended location [55, 112, 113] and for locations containing distracting elements [50].

Activation patterns in the present study show that the BOLD changes occur in the expected parts of the visual cortex. We find activation in anterior regions of the contralateral occipital cortex, which corresponds to the location of the attended checkerboard. Since the checkerboard crosses the horizontal meridian, both the dorsal and ventral parts of V1-V3 are activated. Though V1 activation is often found on an individual level it is relatively weak and the effect is washed out in the group analysis, see Figure 2.8. Without a full retinotopic mapping we can not know for certain which visual areas correspond to the activation clusters, but the activations close to the posterior part of intraparietal sulcus (Figure 2.8 and Figure 2.9) are likely V3A and/or V3B.

Almost all subjects showed activity at the ipsilateral occipital pole (Figure 2.9), a region representing the foveal part of the visual field. Figure 2.11 informs us that the effect in this region is a deactivation during contralateral attention, i.e. the right occipital pole gets suppressed during left peripheral attention. The same effect was reported by Brefczynski-Lewis et.al. [52].

Though the overall pattern was the same across subjects there were also variations, both in location and size of activation clusters. This is partly due to the fact that the anatomical locations and sizes of the visual field maps vary across individuals [47, 118, 119], but on top of this there is also an individual variation in the attentional

topography, e.g. amount of ipsilateral effect and the spread of the activation [52]. However, the individual pattern is consistent and does not change over sessions [52] which is important when considering BCI and implantation of electrodes.

Control signals and classification

The fixed thresholds used for classification in subjects 1-7 turned out to be very conservative, resulting in most images being classified as 'off'. For those subjects the average TPR were 0.41 and 0.25 for left and right, respectively. The adaptive thresholding applied to the two other subjects greatly improved the online sensitivity, while still limiting the false positives. With this improvement, these subjects' TPR averaged 0.75, for left, and 0.85, for right. This motivated an offline re-computation of the first group's performance using the same adaptive method, increasing the TPR to 0.86 and 0.75 for left and right, respectively.

Besides these numbers, based on individual images, we computed a measure of performance by classifying each complete trial. However, since the aim was to test the stability of our control paradigm and its capacity in the context of implanted electrodes, not to optimize the BOLD classification, we avoided time averaging of the data. Instead we also classified each trial using only the 5th time point. This still gave an average correct classification of 89% for left attention and 88% for right attention (Figure 2.7).

These numbers are in the upper range of what has been reported with EEG based systems using e.g. motor imagery and SSVEP [27, 120]. It should be noted that we have included an 'off' class (central attention), which in practice makes it a three-class paradigm. The inclusion of a 'no-choice' option is something that is often overlooked in BCI studies [121]. If we would have classified each time point using only the options of left or right attention, the performance would have been even higher.

Classification of fMRI data is inevitably slow since the BOLD response has a delay of around 5 seconds after neural firing, and it takes a long time before the signal returns to baseline. However, the time delay will not be present in a true BCI system based on electrophysiological signals. Naturally, a quick detection is desired also for our purpose of task evaluation and subject training, but here the few seconds delay is more acceptable.

Suggested improvements

The thresholds should be estimated online as was done for two of our subjects. In this way one can take advantage of the individual differences. The ROI selection can be improved in several ways. As a starting point, we used a fixed number of 500 voxels to include in each ROI. However, the number of voxels selected should probably not be a fixed value but somehow depend on the t-value distribution. On the other hand, a fixed t-value threshold could lead to unpredictable results due to a large variation of the ROI sizes across subjects. It would also be possible to put anatomical restrictions on the ROIs. By defining a mask based on a structural image the voxel selection can be restricted to e.g. a single hemisphere.

Potential

We have shown that the BOLD response following a covert shift of attention to a peripheral region in the visual space is strong enough to be classified in a single trial. Although BOLD is an indirect measure of neural activity, the spatial locations identified by fMRI have been shown to closely match those found using invasive electrophysiological measurements [43, 89, 90]. Despite a limited number of trials and the non-optimal placement of the electrode grid, our ECoG data show that it possible to classify the same attention task using the power in the gamma band. Hence, it is likely that signals recorded by electrodes placed at the optimal positions, as located by fMRI, can be classified with at least the accuracy of our fMRI system. Moreover, the detection will be much quicker based on the electrical response, compared to when using the hemodynamic response.

The spatial attention strategy has some attractive features not found in the tasks commonly used for BCI, such as motor imagery. First, the degrees of freedom can be increased by simply adding more peripheral target regions. Second, a target region can be moved to the location in the visual field that is mapped to, and activates, the cortical area most suitable for implantation. It is also possible that by using this property, and selecting to activate a superficial brain area, it will be easier to pick up the signal changes with EEG or fNIRS.

The real-time fMRI setup described here can be used for evaluating new paradigms as potential control tasks, and to train subjects in them. When planning implantation of intracranial electrodes, the BCI setup can be tried out before surgery in order to locate the best and most stable positions.


3

Real-Time Decoding of Visuospatial Attention to Four Directions

Based on: P. Andersson, N.F. Ramsey, M. Raemaekers, M.A. Viergever and J.P.W. Pluim, "Real-Time Decoding of Direction of Covert Visuospatial Attention", Journal of Neural Engineering, (In Press)

Abstract

Brain-Computer-Interfaces (BCI) make it possible to translate a person's intentions into actions without depending on the muscular system. Brain activity is measured and classified into commands, hereby creating a direct link between the mind and the environment, enabling e.g. cursor control or navigation of a wheelchair or robot. Most BCI research is conducted with scalp EEG but recent developments move towards intracranial electrodes for paralyzed people. The vast majority of BCI studies focus on the motor system as the appropriate target for recording and decoding movement intentions. However, properties of the visual system may make the visual system an attractive and intuitive alternative. We report on a study investigating feasibility of decoding covert visuospatial attention in real-time, exploiting the full potential of a 7T MRI scanner to obtain the necessary signal quality, capitalizing on earlier fMRI studies indicating that covert visuospatial attention changes activity in the visual areas that respond to stimuli presented in the attended area of the visual field. Healthy volunteers were instructed to shift their attention from the center of the screen to one of four static targets in the periphery, without moving their eyes from the center. During the first part of the fMRI-run the relevant brain regions were located using incremental statistical analysis. During the second part the activity in these regions was extracted and classified, and the subject was given visual feedback of the result. Performance was assessed as the number of trials where the real-time classifier correctly identified the direction of attention. On average, 80% of trials was correctly classified (chance level $< 25\%$) based on a single image volume, indicating very high decoding performance. While we restricted the experiment to five attention target regions (four peripheral and one central), the number of directions can be higher provided the brain activity patterns can be distinguished. In summary, the visual system promises to be an effective target for BCI control.

3.1 Introduction

Any loss of control over our muscular system restricts our ability to interact with other people as well as with our physical environment. In severe cases, complete paralysis of all voluntary muscles fully detaches the mind from the environment, a condition referred to as 'locked-in syndrome' [2]. With Brain-Computer-Interfaces (BCI) this broken link between intention and action is reestablished by creating a direct communication channel between a person's brain and the external physical world, without depending on the muscular system [92]. This is done by measuring cortical activation and translating detected changes into commands. For this to work, the user needs one or more mental control tasks to perform in order to regulate local brain activity.

A common target for BCI control is the motor system, where the user imagines moving e.g. a hand, an act that gives a response similar to an actual movement [38, 40, 100, 122]. Other common BCI control strategies are the P300 response [22–24] and steady state visual evoked potentials (SSVEP)[16–20]. The P300 EEG response is evoked by a rare or surprising event. A P300 BCI has a computer screen displaying a number of choices that flash in a random sequence. When the user focuses on one of the choices, a P300 response will be elicited when this particular choice lights up. SSVEPs are induced by flickering visual stimuli to which the users either move their gaze, or focus their attention. By using different frequencies for different options one can link the SSVEPs to a particular choice of action. It is worth noting that both P300 and SSVEP based BCIs are intrinsically dependent on external visual stimulation. Though BCI control has been accomplished using motor imagery, P300 and SSVEP, these strategies were originally chosen because they are accessible by scalp electrodes.

There is now a growing interest in methods where electrodes are implanted, either on the dural or cortical surface using electrocorticography (ECoG), or directly into the gray matter [38, 40, 43, 97–99]. Moving to intracranial modalities is motivated by studies where successful BCI control is established in non-human primates [36, 95, 96]. Intracranial electrodes give access to stronger signals at higher spatial resolution and with a higher bandwidth compared with scalp electrodes. With increased signal strength it is hoped that a more stable and reliable BCI control can be achieved. The higher resolution has the potential of distinguishing more "brain states" meaning that a more advanced control can be created. Intracranial recordings can also access deeper brain regions and possibly reach functional systems previously inaccessible. Thus, besides providing a better signal, intracranial recordings can potentially open up for novel brain functions for BCI control.

Since BOLD signal changes measured with fMRI spatially match direct electrical

signals measured with e.g. ECoG [43, 89–91], it is the best noninvasive tool in the preparation and planning for BCI implants. fMRI can for instance help determine patient-specific details regarding the optimal implantation site and brain function for BCI control. Moreover, using real-time fMRI [9] with online feedback, it is possible for the patient to practice BCI control before any surgical procedures.

An intriguing candidate for BCI control is the visual system. Each position in the visual field is represented by a specific location in multiple cortical field maps. This mapping has been well studied using fMRI, see e.g. [47]. Visuospatial attention shares important brain activity features with actual visual stimulation. Importantly, merely directing attention covertly (hence maintaining gaze at the center) to a particular position in the visual field induces local changes in neural activity that are highly similar to activity in response to actual visual stimulation in the same position [48–55]. It should be noted that this effect is not dependent on eye movements or visual input, as is the case with SSVEP and P300. Covert directing of visuospatial attention can be performed quite well, i.e. without moving the eyes, as has been shown in many studies [49, 51, 106, 123, 124]. We have shown in Chapter 2 that two directions of covert visuospatial attention can be decoded in real time from fMRI data. The possibility of using stimuli independent covert visuospatial attention to control a BCI has also been explored using offline analysis of changes in the alpha frequency band in EEG [107, 125] and MEG [126, 127]. So far, however, real-time decoding of two directions (left and right field) has not yet been successful with scalp EEG (nor with MEG), leaving the question of feasibility unanswered. Most covert attention studies report on offline analysis, which is a step away from real proof of principle since it allows for optimizing feature selection and classification on existing data. In the present study we investigate feasibility of decoding direction of covert attention in four directions, in real-time, from occipital and parietal cortex. To be able to do this, we utilized the exceptional contrast to noise ratio of a 7 Tesla human MRI scanner [69, 128].

3.2 Material and methods

3.2.1 Subjects

Nine healthy volunteers (age 23-28, right-handed, 5 male) with normal or corrected-to-normal vision participated in the study, after giving their written informed consent. All subjects were naive to the task. The study was approved by the ethics committee of the University Medical Center Utrecht in accordance with the declaration of Helsinki

(2008). Activated regions were first localized in real-time using incremental statistical analysis. This was directly followed by a second part where the subjects received feedback based on the activity in these regions.

For three of the subjects additional fMRI data were acquired with the purpose of mapping their retinotopy [47]. With these data we could relate the selected feedback regions to known visual field maps in the occipital cortex.

3.2.2 Data Acquisition

Real-time fMRI was implemented using a 7T Philips Achieva system with a 16-channel headcoil. The functional data were recorded using an EPI sequence (TR/TE= 1620/25ms; FA=90; SENSE factor=2; 35 coronal slices, acquisition matrix 96x96, slice thickness 2mm with no gap, 1.848 mm in-slice resolution). The FOV was selected so it covered the occipital lobe and the most posterior part of the parietal lobe. Each experiment consisted of 652 volumes. Directly following reconstruction on the scanner, the data were sent to a separate computer performing the analysis (Dual-Core 2.5GHz notebook) via the local network using a TCP/IP protocol and the Philips DRIN (Direct Reconstruction INterface) module (see Figure 3.1a). The stimulus was projected to the subject from a second computer via a video projector. Update-triggers containing information about the new instruction and color of the instruction cue (for feedback) were sent to the second computer via a serial cable.

A high-resolution image was acquired for the anatomy using a T_1 3D TFE sequence (TR/TE=6/2ms; FA=7; FOV=220x180x200mm; 0.55x0.55x0.5mm reconstructed resolution).

For three subjects we acquired additional data for retinotopic mapping using the same EPI sequence and FOV as in the attention experiment. A total of 340 volumes were collected, evenly divided in polar and eccentricity mapping stimuli.

3.2.3 Task

During the experiment the subjects were presented with a screen as seen in Figure 3.1b. The screen contained a central cue on which the subjects were instructed to fixate their gaze at all times. Four yellow arrows indicated the targets for left, right, up and down attention and were all constantly visible. The center cue alternated between a circle and an arrow pointing towards one of the yellow target areas. Subjects were instructed to covertly move the attention to the target area located in the direction

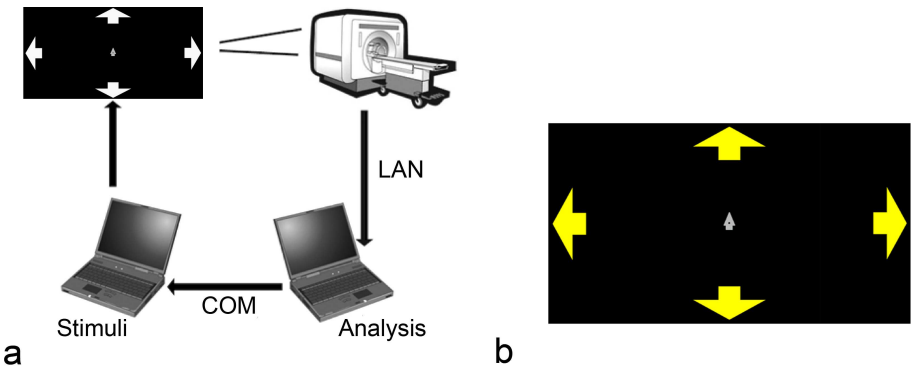

Figure 3.1: (a) The real-time system setup. (b) The screen presented to the subject. While keeping the gaze fixated at the center, the subject shifts the attention to the yellow area in the direction of the central cue.

of the instruction arrow or, in the case of a circle, to relax and keep attending to the center.

A beamer projected the image into the bore of the scanner onto a small projection screen close to the head of the subject. This limited the size of the projected image and it was more restricted vertically than horizontally. We wanted the target regions as peripheral as possible on the screen. This positioning increases the distance between the mapped brain areas, making them easier to separate in the classification. Furthermore, when using the setup in an actual BCI application the foveal region will be free to use for visual feedback etc. As a consequence of the limited image projection size the distance from the central point to the up and the down regions was shorter than the distance to the left and right regions (up/down 6° and left/right 11° angle from the central line).

3.2.4 Motion Correction

The first volume of the fMRI run was saved and used as a template to which all the subsequent volumes were aligned. The rigid registration was performed by minimization of the sum of squared differences between gray-value intensities. A stochastic gradient descent method [111] was employed for optimization, using 50 iterations. Prior to the registration, the images were blurred with a Gaussian filter (σ = 1 voxel). Linear interpolation was used during the optimization steps while the final image was generated using cubic B-spline interpolation. The algorithm was implemented in C++,

and called from Matlab (Mathworks, Natick, MA). The computation time was approximately 0.6s per fMRI volume.

3.2.5 Localizing part

Trials

The first 292 volumes were acquired to locate the activated regions. The four attention conditions (directions) were randomized and repeated nine times. Each trial was 6.48s (4 TRs) long followed by the same length of rest condition (circle cue).

Analysis

The localizing images were statistically analyzed on the fly using the incremental GLM algorithm described in [87]. Five regressors representing the four directions of attention plus a linear drift term were included to compute four t-maps using the contrasts 'one minus the others'. That is, for right attention the contrast was 'right - $\frac{1}{3}$[left+up+down]' etc. The reason for contrasting these differences was that besides an increased BOLD signal in retinotopically mapped regions, the attention can induce a decreased signal in surrounding and unattended regions [50, 55, 112, 113]. A second reason was to avoid selecting regions responding to attention in general.

ROI selection

When the last volume in the localizer part had been analyzed, the resulting t-maps were used to define the ROIs representing the activated voxels. From each t-map the 500 highest values were first selected, and from this selection, clusters smaller than 5 voxels were removed. The resulting masks were used as $\mathrm{ROI_R}$, $\mathrm{ROI_L}$, $\mathrm{ROI_U}$ and $\mathrm{ROI_D}$, for right, left, up and down respectively. For each voxel inside an ROI, the baseline and standard deviation over time were estimated and later used for normalization (Equation 3.1).

3.2.6 Feedback part

Trials

During the next 360 volumes, activation in the selected regions was processed (classified) on the fly and feedback was given to the subject based on the classification

result. Here the trials were 9.72s (6 TRs) long followed by 4.86s (3 TRs) of rest (circle cue). The four directions were repeated in 10 trials each. At the start of a new trial the instruction was not updated until the analysis (motion correction, detrending, classification etc) of the last image volume in the previous trial was finished. As a consequence the start of the trials were delayed by on average 1.1 seconds. This has been accounted for in all plots and results.

BCI Control Signals

The classification of an image volume was done based on four control signals, each representing one of the attention directions. These control signals were in turn based on the average values inside the ROIs determined during the localizing part. When a new volume had been motion corrected, the values inside the ROIs were extracted and detrended using the algorithm described in [114]. For each ROI_d the detrended voxel values were first individually normalized and then averaged to give a single value as

$$\tilde{v}_d = \underset{k}{\text{mean}}\left(\frac{v_d^k - m_d^k}{\sigma_d^k}\right), \text{ for all } k \in \text{ROI}_d \tag{3.1}$$

where m_d^k and σ_d^k are the estimated baseline and standard deviation for voxel k. Finally, these numbers were subtracted, using the same contrasts as for the corresponding t-map, giving four values S_R, S_L, S_U and S_D, e.g. $S_R = \tilde{v}_R - \frac{1}{3}[\tilde{v}_L + \tilde{v}_U + \tilde{v}_D]$.

Classification

An image volume could be classified either as 'active', attention directed to one of the four directions, or 'off', attention directed to the center. Each control signal was defined to increase during one of the 'active' conditions. If none of the S-values were above a fixed threshold the volume was classified as 'off'. This threshold was determined from earlier pilot experiments and set to 0.3. When one or more control signals exceeded the threshold, the volume was assigned to the 'active' class with the highest control signal.

Feedback

Feedback was given to the subject via a change in color of the central instruction symbol. During the localizer part, where no feedback was given, the color was constantly gray. During the feedback part the color was green if the direction classified as active agreed with the instruction, and red if it did not agree. Note that feedback was given

also during the central condition, where a green color was given for volumes classified as 'off'. Before the experiment the subject was informed about the consequences of the hemodynamic delay, and that the effect of the attention shift would only be seen at the end of a trial.

Performance

As an offline measure of performance we computed the percentage of correctly classified trials. Since each trial consisted of six volumes but we are interested in the classification of single volumes, we computed the performance for each of the six time points individually. This also gave us an estimate of the strength of the control signals over time and how quickly the attention shifts can be detected.

3.2.7 Group maps

A group map was computed in order to locate the most commonly selected regions for ROIs over subjects. Each subject's structural T1 image were first normalized to the Montreal Neurological Institute (MNI) space using SPM5. Using the result from this spatial registration, the four ROIs were transformed to MNI space. With all subjects' ROIs in the same space they were now added. The four resulting maps show, with a number between zero and nine, for how many subjects a voxel was included in the corresponding ROI. Though there are inter-subject variations in size and location of the visual field maps [47, 118, 119], these group maps will give an indication of the most commonly activated regions.

3.2.8 Retinotopic mapping

For three of the subjects (Subjects 7-9) two additional scans (polar and eccentricity) were acquired to locate the visual field maps V1/V2/V3. Since visuospatial attention follows the same mapping as stimulation, the locations of the ROIs could be directly compared to these retinotopic maps. A directed gaze, instead of directed attention, would give a distinctively different activation pattern. The location of the ROIs within the visual field maps can therefore validate that our findings are based on covert attention and not eye-movements.

The mapping was done using a rotating wedge stimulus for the polar angle mapping, and concentric expanding rings for eccentricity mapping. Both the rings and the wedges consisted of a flickering (8 Hz) radial checkerboard where the check size, and

ring thickness, varied with eccentricity to compensate for the cortical magnification factor. The angle of the wedge was 45^o. The rotation and the eccentricity cycles each took 55 s (34 TRs) to finish and 4 cycles were acquired.

The analysis was done as described in e.g. [129, 130], i.e. by computing the phase of the periodic response at the rotation frequency.

3.2.9 Flatmapping

To better visualize the locations of the ROIs within the retinotopic maps, the occipital lobes were inflated and flattened using the software Caret [131]. For the flattening, a cut was made in the cortical surface along the calcarine sulcus.

3.3 Results

3.3.1 Performance

Figure 3.2 gives an overview of the classification results. It displays the number of subjects with a correct classification for every volume in each of the trials. To visualize the hemodynamic delay's effect on the classification, the performance was computed for the six time points (volumes) within the trials. The plotted curves in Figure 3.2 show the percentage of the different time points that on average could be correctly classified. (Note that the time is shifted 1.1 seconds according to the time of instruction onset (see 3.2.6).) In other words, the plots show the percentage of trials correctly classified when using a single image volume acquired at different times after the attention shift.

The performance curves in Figure 3.2 can be compared to the plots in Figure 3.3, showing the response in the control signals during the four trial types, averaged over all subjects. The averages were computed over both the localizer trials and the feedback trials and therefore only span the 4 TRs of the shorter localizer trials. Note that these are not the raw BOLD responses, but the contrasts between the responses in the four ROIs.

When averaged over both subjects and attention directions the percentage of correctly classified trials based on the fifth volume was 79.4% (Table 3.1). The individual subjects' averages varied between 63% and 100%. Three subjects had a single direction with accuracy lower than 50%.

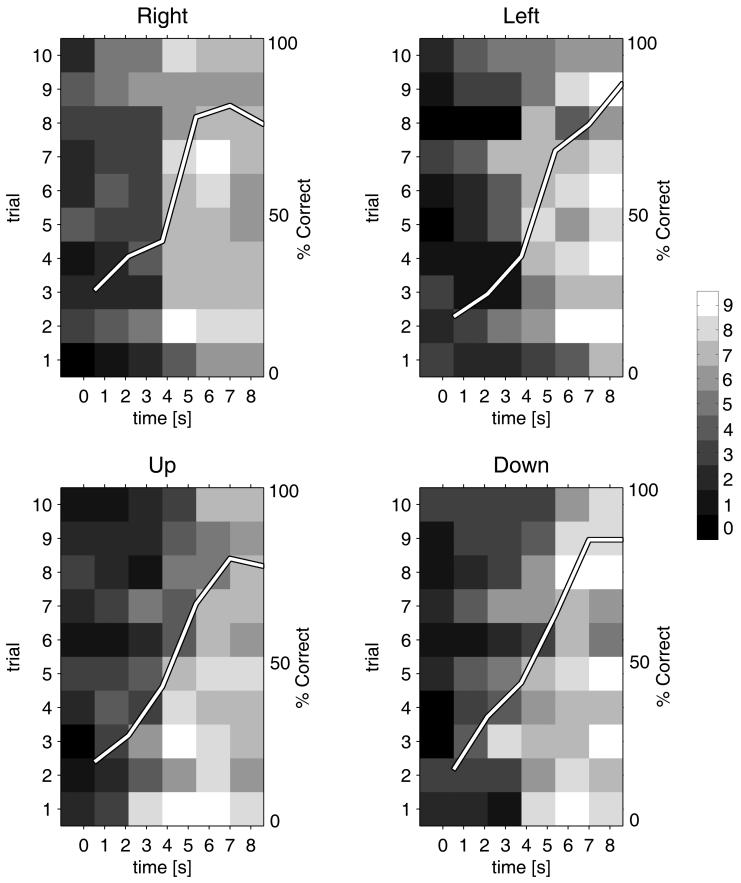

Figure 3.2: Number of subjects having a specific image volume correctly classified. Each row represents one of the 10 trials, and each column one image volume in that trial. The color of each square shows how many subjects had this volume classified correctly. The curves show, for all volumes, what percentage of the trials that would be correctly classified if based only on this particular volume. The scale of the horizontal axis is in seconds and the time has been shifted 1.1 seconds according to the time of instruction onset (see 3.2.6)

3.3.2 Group maps

Figure 3.4 shows the group maps computed by adding the spatially normalized ROIs. The range of the color scale is from zero to nine and represents the selection frequency. To emphasize the most important regions the maps are thresholded at unit value in Figure 3.4, i.e they only show regions selected for more than one subject.

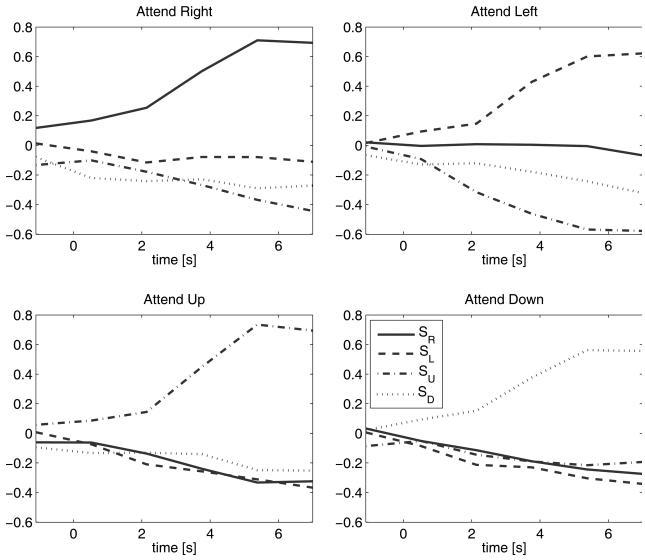

Figure 3.3: The average control signals, S_R, S_L, S_U and S_D during the four conditions.

Table 3.1: Percentage of the ten trials for each attention direction that were classified correctly, if based only on the fifth volume in each trial. (R,L,U,D = right,left,up,down attention)

	Subject									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**Avg**
R	80	90	100	90	90	80	90	50	50	80.0
L	100	60	100	80	80	80	20	80	70	74.4
U	90	30	100	30	90	100	70	100	100	78.9
D	70	100	100	50	100	80	90	70	100	84.4
Avg	85.0	70.0	100.0	62.5	90.0	85.0	67.5	75.0	80.0	79.4

3.3.3 Retinotopic mapping

The organization of the human visual cortex is well studied [47]. In V1, and in the encircling areas V2 and V3, both hemifields map onto the contralateral hemisphere, and stimuli above and below the horizontal meridian are mapped ventral and dorsal to the calcarine sulcus respectively. From the center of the visual field to the periphery, cortical representations are laid out from the occipital pole towards more anterior aspects of the visual cortex. Using Subject 9 as an example, Figure 3.5 shows how the horizontal and vertical meridians are represented in V1-V3. The borders between V1, V2 and V3 together with the meridians' representations are displayed on the inflated

brain. Several other visual maps exist that process spatial information [47]. Adjacent to the anterior part of dorsal V3 are two full hemifield maps, V3A and V3B, and additional smaller maps are distributed towards the posterior parietal lobes. All subjects' ROIs have voxels that are located in dorsal regions outside V1-V3. However, to find what visual field maps these locations correspond to we would need longer retinotopic mapping sessions and a more sophisticated analysis.

Figure 3.6 (top) shows the locations that in the polar mapping showed the strongest response for the phases corresponding to the attention target areas ($0, 90, 180$ and 270°). Figure 3.6 $(bottom)$ displays, on the flattened surfaces, how the ROIs were distributed within V1, V2 and V3 (borders shown in black) for Subject 9. The targets for left and right attention cross the horizontal meridian, and thus as expected (Figure 3.5) voxels are selected in both the dorsal and ventral parts of V1-V3. Similarly, since the target regions for up and down cross the vertical meridian, voxels are selected in both the left and the right hemispheres.

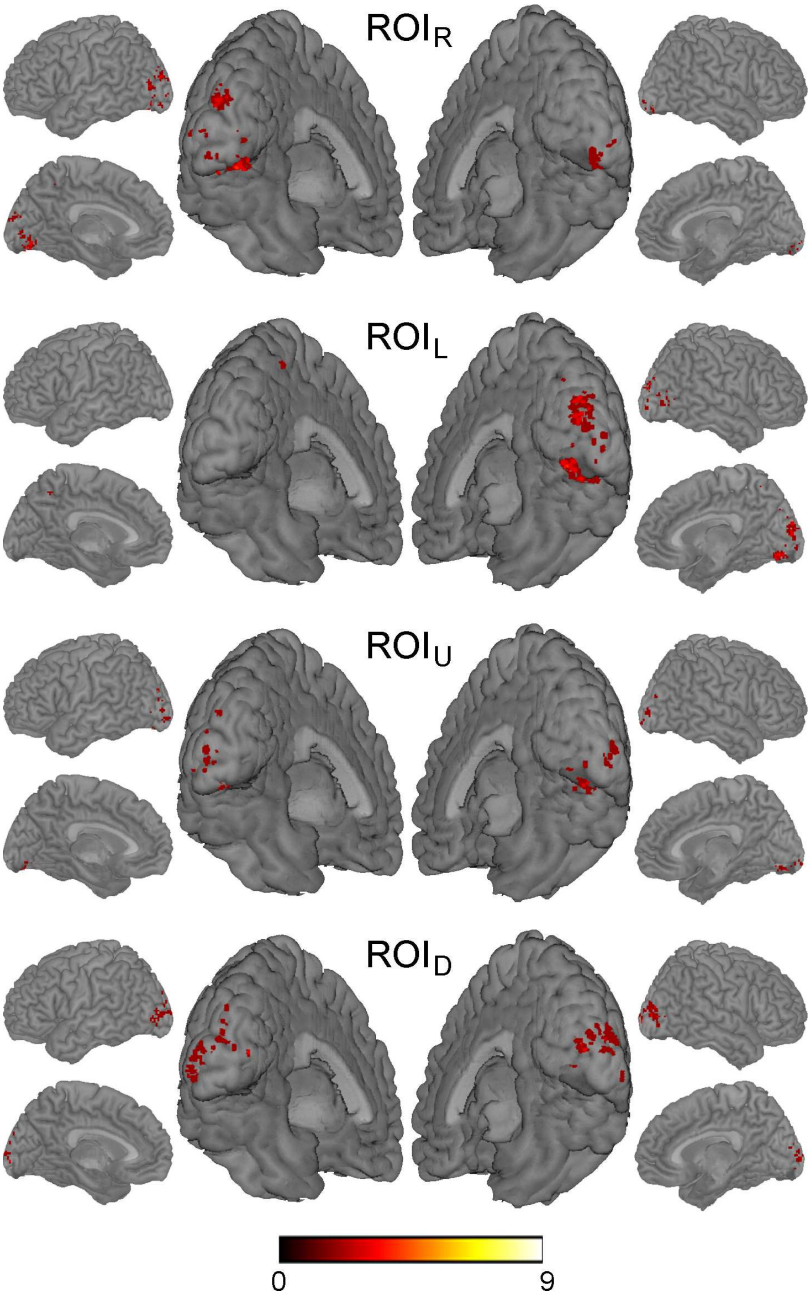

Figure 3.4: Group maps computed from all subjects' ROIs, overlayed on the MNI brain. The colorscale represents the number of subjects for which the voxel was selected. The voxels shown are the ones with values ≥ 2, i.e. voxels that were included in two or more subjects' ROIs.

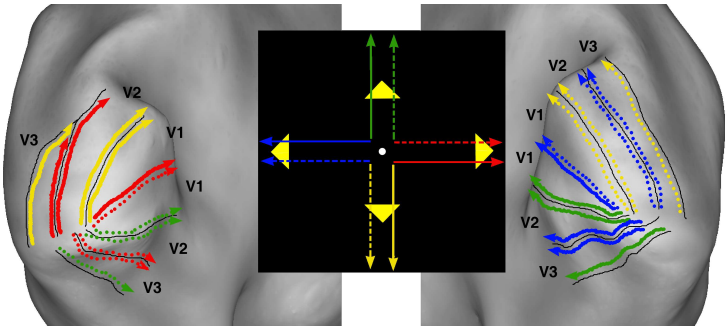

Figure 3.5: The black lines show the borders between V1, V2 and V3 (Subject 9) and the red, blue, green and yellow lines show how the horizontal and vertical meridians are represented in these areas.

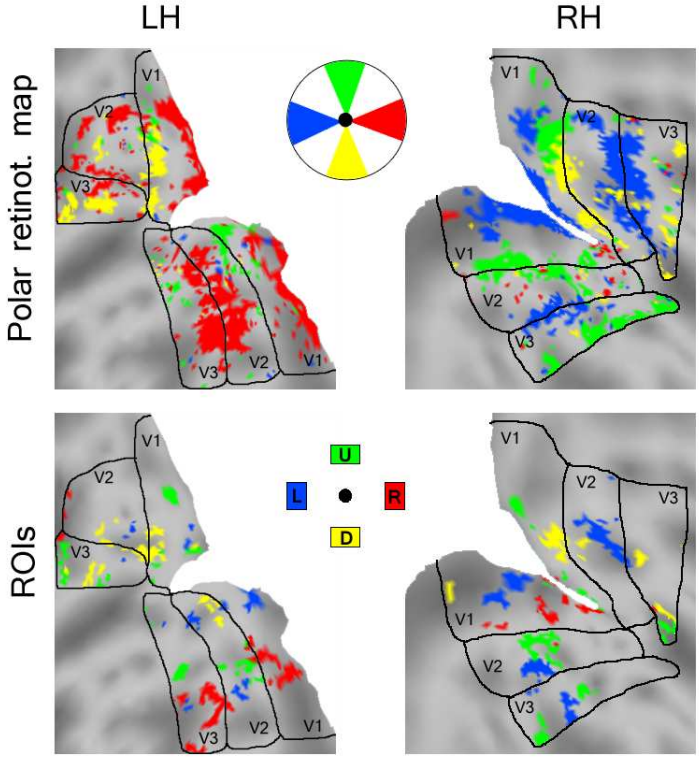

Figure 3.6: The ROIs for Subject 9 projected on the flattened cortical surfaces. Only the parts within V1-V3 are shown (borders shown as black lines). (top) : The locations that in the polar mapping showed the strongest response for the phases corresponding to the attention target areas. $(bottom)$: The locations of $\mathrm{ROI_R}$, $\mathrm{ROI_L}$, $\mathrm{ROI_U}$ and $\mathrm{ROI_D}$.

3.4 Discussion

The results presented here show that intentional covert directing of visuospatial attention to one of multiple targets can be detected and classified in real-time. This effect is fully independent of any visual stimulation and eye movements. As a measure of performance we classified each trial based on the fifth image volume in the trial. This gave an average of 79.4% correct classifications, which is clearly significant ($p < 10^{-16}$; binomial test, assuming symmetric class distributions, chance level 0.25). Since all out test trials are from one of the four peripheral directions but we classify into five classes where the centre condition is included, the level of chance is less than 25%. The individual performances varied between 63% and 100%. It is notable that whereas three subjects had a single direction with a hit-rate lower than 50%, all subjects had at least two directions with 80% or more. The reason for the intra-subject differences in performance for the different directions is unknown. All subjects had normal or corrected-to-normal vision. While the correct classification rates are much higher than the level of chance, the performance would still most likely improve if the subjects were given a chance to practice the task.

The three subjects for which a retinotopic mapping was performed show that the voxels selected for feedback are mainly in the expected parts of V1-V3. Also, voxels located along more dorsal areas towards the parietal lobe are commonly selected. This is also expected since this part of the cortex contain several functional areas processing visuospatial information, e.g. V3A/B.

The high field prohibited the use of an eye tracker for safety reasons. However, a large number of studies (e.g. [49, 51, 106, 123, 124]) have shown that people can perform covert shifts of visuospatial attention without moving the eyes. Moreover, the retinotopic maps acquired in 3 subjects (Figure 3.6) and the group maps (Figure 3.4) confirm that we decoded covert attention and not eye-movements. The pattern of activation as revealed by the ROIs selected on the basis of directional contrasts closely match those obtained with the retinotopic mapping procedure.

If the subjects were to move the eyes towards the target, the resulting activation patterns would be distinctly different, most likely appearing mirrored (up-down and left-right). When the gaze moves towards one of the targets, in retinotopic space, all targets move towards the opposite hemifield. For instance, if the gaze is moved to the left, all four targets are moved to the right in the visual field. The up and down targets will be positioned completely in the right hemifield. As a result, the largest effect of a movement to the left would be seen in the ipsilateral left hemisphere. Further, since the visual stimulation is constant, an eye-movement would only result

in a transient response that, if detectable, would not remain as long as the responses seen in Figure 3.3.

The BOLD response is very slow compared to a direct electrophysiological measure. However, a short delay in the feedback is not a problem when the purpose of the real-time fMRI is to practice or evaluate a brain function for BCI control.

The low spatial resolutions of EEG and MEG suggest it will be difficult using those methods to separate activation patterns from multiple attention target regions. Furthermore, there is a large variability over subjects in the distribution of the retinotopic field maps. This will make the cortical representation of a particular part of the visual field differ in distance from the scalp, as well as in its orientation to the scalp (affecting MEG). van Gerven et al. [126] used MEG and the power in the alpha band to decode covert visuospatial attention to four peripheral areas, using a setup similar to ours. The study showed an average correct classification of 41% compared to our 79%, even though they did not classify the central condition. Since they classified electrophysiological signals, not affected by a hemodynamic delay, they could however use a shorter attention interval of 2.5 seconds.

By using ECoG with its high spatial resolution and signal strength it would likely be possible to reliably separate multiple directions of attention. With ECoG it would also be possible to use the higher frequencies in the gamma band. These frequencies have been shown to have the strongest spatial correlation to BOLD [91].

In conclusion, we have shown that covert visuospatial attention towards multiple target regions can be decoded in real time using fMRI. All our subjects could reliably perform the task without any training. Visuospatial attention based paradigms should therefore be considered an option for multi-dimensional BCI control. It will be of particular interest for BCIs using intracranial electrodes, since those will be able to detect and decode the local activation changes from multiple directions of attention.


4

Navigation of a Telepresence Robot via Covert Visuospatial Attention

Abstract

Brain-computer interfaces (BCIs) allow people with severe neurological impairment and without ability to control their muscles to gain back some control over their environment. The BCI user apply a mental task to regulate brain activity that is measured and translated into commands controlling some external device. We have investigated whether people are capable of navigating a robot through a track by shifting their visuospatial attention. Visuospatial attention constitutes a very intuitive brain function for spatial navigation and the control does not depend on presented stimuli or eye movements. Our robot is equipped with motors and a camera that sends visual feedback to the user who can navigate it from a remote location. We used an ultrahigh field MRI scanner (7 Tesla) to obtain fMRI signals that were decoded in real time using a support vector machine. The results show that BCI based on visuospatial attention can let paralyzed people without any motor control navigate around and communicate with their environment, even from a remote location.

4.1 Introduction

The concept of Brain-Computer Interfaces (BCI) concerns technologies creating direct communication channels between the brain and a computer or other type of device. The goal is to accomplish real-time decoding of brain activity with sufficient reliability for paralyzed people to use it in their daily life. Two essential and defining components in a BCI system are the modality used for measuring brain activity and the mental control tasks used for regulating this activity. When it comes to the measuring, the main focus has so far been on electroencephalography (EEG). However, the implicit disadvantages of EEG, such as a low spatial resolution and sensitivity to non-neural electrical activity, have created a growing interest in BCI using invasive measuring techniques. By implanting intra-cranial electrodes the quality, bandwidth and spatial resolution of the signal can be increased significantly.
The mental control tasks have until now mainly been based on brain functions that involve strong signals, such as the motor potential and the P300 oddball response [92], since these can be detected well from the scalp using EEG. When considering invasive technologies the improvements in signal quality open up for the use of brain functions previously not tested for BCI use [43, 45, 46]. Since BCI users might find some brain functions harder, or even impossible to control, and some brain functions might be more intuitive for certain BCI applications, an effort should be made to evaluate new BCI control paradigms. The two previous chapters indicate the potential of a new approach using top-down regulation of the sensory cortices via attention (see also [46]). Attention can change brain activity even in the absence of exogenous stimuli [132]. Attending to a region of the peripheral visual field, while keeping the gaze fixed, generates neural responses in the parts of the cortex processing visual information in this region [51, 52]. In Chapters 2 and 3 we showed that it is possible to decode individual fMRI images in real time during covert visuospatial attention. When studying the feasibility of a BCI control strategy it is important that it is tested in a closed-loop setup. It is only when the test subject gets to adapt to a real-time feedback, preferably within a real-life application, that the control strategy can really be evaluated.

Visuospatial attention would constitute a very intuitive brain function for spatial navigation. In the present study we therefore investigate whether people are capable of navigating a robot by shifting their visuospatial attention, without the aid of presented stimuli or eye movements. The subjects were instructed to navigate the robot through a track containing targets that should be reached in a particular order. Our robot is equipped with a camera and the image is sent as feedback to the user. We used an

ultrahigh field MRI scanner (7 Tesla) to obtain an fMRI signal that is strong enough for real-time decoding. Since BOLD activity is well correlated spatially with changes in the higher frequencies of electrophysiological signals [89, 91], the performance with fMRI is a good indication of the feasibility of a BCI with electrode implants.

4.2 Materials and Methods

4.2.1 Subjects

Four healthy volunteers (age 20-50, right-handed, 2 male) with normal or corrected-to-normal vision participated in the study, after giving their written informed consent. The study was approved by the ethics committee of the University Medical Center Utrecht in accordance with the declaration of Helsinki (2008). Each subject was scanned three times (one practice session and two performance sessions) separated by between 1 and 28 days.

4.2.2 Robot

We used the Erector Spykee robot (Meccano Toys Ltd) that is equipped with a wireless modem and a video camera. The software was made so that a forward movement instruction moved the robot approximately 50 cm, while a right or left instruction turned it approximately 30 degrees. The robot had no mechanism preventing it from hitting the wall. When on a few occasions the robot moved forward and ended up with the front flat against the wall, it was moved back to the previous position and orientation.

4.2.3 Data

The subjects were scanned at a 7T Philips Achieva system with a 16-channel headcoil, which generates the signal quality needed for our purpose. The functional data were recorded using an EPI sequence (TR/TE=1620/25ms; FA=90; SENSE factor=2; 35 coronal slices, acquisition matrix 96x96, slice thickness 2mm with no gap, 1.848 mm in-slice resolution). The field of view (FOV) was selected such that it covered the occipital lobe and the most posterior part of the parietal lobe. A high-resolution image was acquired for the anatomy using a T_1 3D TFE sequence (TR/TE=6/2ms; FA=7; FOV=220x180x200mm; 0.55x0.55x0.5mm reconstructed resolution).

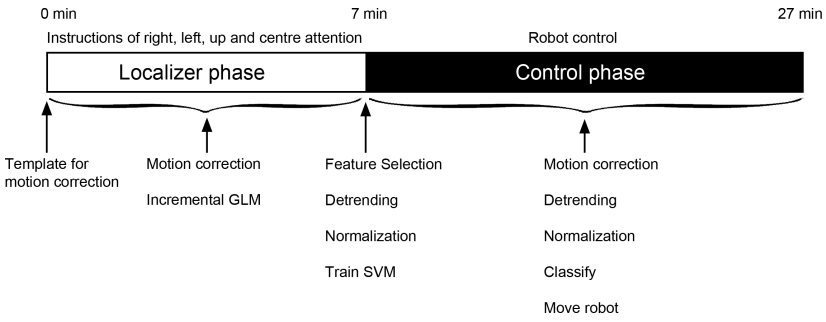

Figure 4.1: Illustration of the experiment timeline.

4.2.4 Experimental Setup

Each session consisted of a single fMRI run of 995 image volumes. The first 270 volumes, the localizer phase, were used for locating relevant voxels and training the classifier whereas the remaining 725, the control phase, were classified as commands to control the robot. During the localizer phase the subjects were instructed where to attend. Trials of right, left and up attention were randomized, and always separated by a center attention trial. Each trial was 8.1s (5 TRs) long. During the control phase there were no instructions given and the subject could move the attention at will. Up attention now made the robot move forward while left and right attention resulted in a turn to the respective direction. The gaze were fixated during the complete experiment and there were no eye movements involved in the control. A timeline of the experiment, showing the different steps of the online analysis, can be found in Figure 4.1.

4.2.5 Task and Navigation Interface

During the experiment the subjects were presented with an image as in Figure 4.2, projected onto a small projection screen in the bore of the scanner. A circle, on which the subjects were instructed to fixate their gaze at all times, was displayed at the center of the screen. Three yellow triangles permanently positioned to the sides of and above the central area indicated the attention target areas, used for sending commands.

During the localizer phase (Figure 4.2a), the center cue alternated between showing an arrow pointing towards one of the yellow target areas and showing only the circle. Subjects were instructed to covertly move the attention to the target area lo-

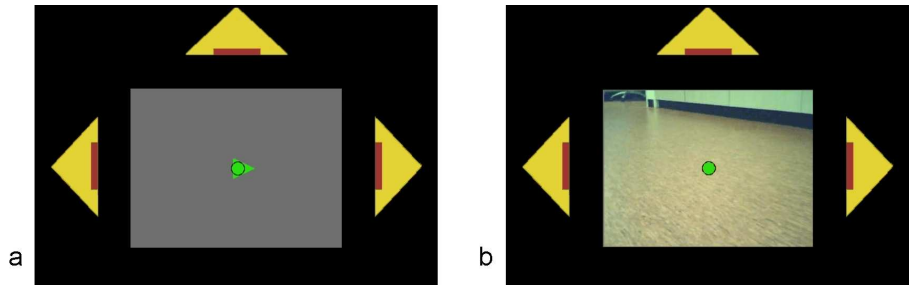

Figure 4.2: The feedback screen. The screen projected to the user during (a) the localizer phase, and (b) the control phase. The three yellow triangles served as targets for left, right and up attention. The green circle in the center indicates the point where the gaze was to be focused at all times. During the localizer phase the subjects' attention should shift according to a central cue ((a) shows the cue for right attention). During the control phase the video from the robot's camera was displayed in the central area.

cated in the direction of the instruction arrow or, in the case of a circle, to keep the attention in the center.

During the control phase the live video images were displayed in the area between the attention targets (Figure 4.2b) and there was no instruction cue indicating where to focus the attention. After the execution of a movement, and before the next command was sent, the attention needed to come back to the center for a short time in order to let the hemodynamic effect wash out. To facilitate this, after a volume had been classified as either right, left or up attention and the corresponding movement had been executed, the video was turned off during four volumes (6.48s). Though the hemodynamic response takes longer than that to completely disappear, the BOLD signal has stabilized enough for a new command to be sent.

4.2.6 BCI hardware

The BCI system consisted of two computers communicating in real time with each other, the MR scanner and the robot (see Figure 4.3). One computer received the images from the scanner directly after reconstruction via the local network using a TCP/IP protocol and the Philips DRIN (Direct Reconstruction INterface) module. This computer performed the main analysis (motion correction, detrending, SVM training, classification etc). The second computer contained the graphic display, projected to the subject via a video projector. The display was updated according to instructions from the first computer via a serial cable. The second computer also contained the wireless link to the robot for communicating the video images and the

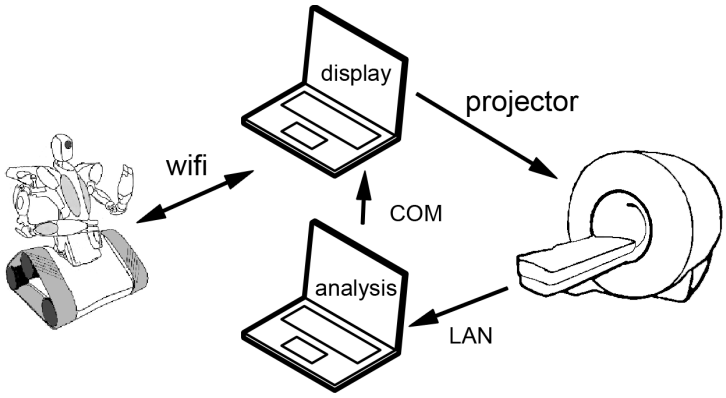

Figure 4.3: BCI system. The setup contains the MR scanner, two computers and the robot.

movement commands. The graphic display and robot communication were implemented using the RoboRealm software (www.roborealm.com).

4.2.7 Motion Correction

All image volumes were corrected for head movements. Motion during the localizer phase will result in a weaker classifier, and during the control phase the wrong features will be extracted from the image volume and sent as input to the classifier. Every image volume was rigidly registered to the first localizer image. Before the fitting, the images were blurred with a Gaussian filter (σ = 1 voxel). As similarity metric we used the sum of squared differences and the optimization scheme consisted of 50 iterations of the stochastic gradient descent method described in [111]. The final image was generated using cubic B-spline interpolation. The code was implemented in C++ and was compiled to a Matlab (Mathworks, Natick, MA) mex-file.

4.2.8 Feature selection

Each sample of fMRI data, i.e. each volume, contains a very large number of voxels of which the majority are either outside the brain or not processing the attention task. In order to avoid overfitting the classifier model, a feature selection step is necessary before it is built. Overfitting occurs when the classifier is trained on voxels that contribute with information irrelevant for determining the attention state.

Our voxel selection is based on a GLM analysis that runs during the localizing part. Four statistical t-maps were incrementally updated with every new image using the algorithm described in [87]. The GLM model contained five regressors; right, left, up and center attention plus a linear drift term. The t-values were computed using the contrasts 'one minus the others'. That is, for right attention the contrast was 'right - $\frac{1}{3}$[left+up+center]' etc. After the last iteration of updating the t-maps, the voxel selection was performed in two steps. A first selection was made by merging the voxels with the 500 highest values from each of the four t-maps. Second, from this first selection clusters smaller than 5 voxels were removed. The remaining pool of voxels determined the ones available for the SVM to train on.

It is possible that a multivariate method such as Recursive Feature Elimination [85], using the actual classification model, could provide a voxel selection giving slightly better performance. However, the computation would take much longer and we would not be able to keep both the localizer and control part in a single fMRI run.

4.2.9 SVM classifier

We used the LIBSVM [133] implementation of a C-SVM classifier with a linear kernel and the regularization parameter $C = 1$. Theoretically, if C is too large, we risk overfitting, and if it is too small, underfitting. However, it has been shown that the classification result is rather insensitive to the value of C [134], and the unit value is often used. LIBSVM uses the "one-against-one" approach for multiclass problems. This means that our classifier consisted of six binary SVMs, one for each pair of classes (attention directions), and an image was assigned to the class with the majority vote. In case of a tie we classified it as center attention (i.e. no action was taken by the robot).

4.2.10 Detrending and Normalization

fMRI signals always contain low-frequency drift to various degrees. To minimize the influence of these signal changes on the classification we applied detrending to the data. For this we used an implementation of the algorithm described in [114] with regularization parameter $\lambda = 200$.

As soon as the last image volume of the localizer phase had been analyzed and the feature selection was ready, the complete time series of the selected voxels were detrended. From the detrended data we then estimated the baseline and standard de-

viation for each voxel. Using these estimates the data was normalized so that each voxel's time series was estimated to have zero mean and unit variance. The detrended and normalized data was then finally used for training the SVM. The original non-detrended values were kept in memory so that they could be used in the detrending of later image volumes. During the control phase, as soon as a new image volume had been passed on from the MRI scanner and had been registered to the template image, the values were detrended and normalized in the same way as the training data. The processed values were then classified by the SVM.

4.2.11 Practice session

The purpose of the first session was for the subjects to get acquainted with the robot control environment and the delay induced by the BOLD response. They were asked to try different lengths of attention to find what worked best for them.

4.2.12 Evaluation sessions

In the two evaluation sessions the robot was placed in the same room as in the practice session. Four targets (25x50cm), marked out on the floor and labeled with the numbers 1-4, were distributed in the room as seen in Figure 4.4. The instructions were to move the robot to these targets in order, and if target four was reached, to continue until the time was up. The time of each target reached was recorded in order to be used as a measure of performance.

4.3 Results

4.3.1 Feature selection

In the feature selection we merged 500 voxels from all four t-maps. However, due to partial overlaps and the second step where we remove any cluster of voxels smaller than five, the final selection consisted of fewer than 2000 voxels. Table 4.1 shows the number of voxels selected and used in the SVM training in each of the sessions. The average number of voxels included was 1236, which corresponds to a volume of 8.4cm^3.

Figure 4.5 shows a group map of the voxels selected from each attention direction (before they were merged to a single selection). Each subject's data was spatially nor-

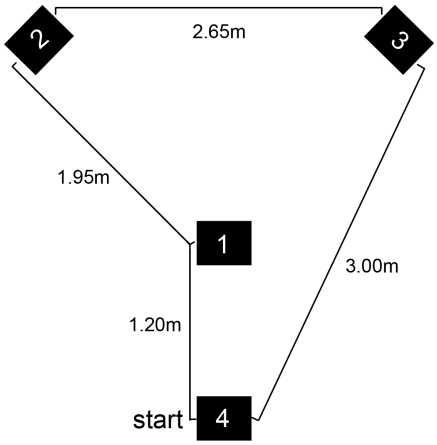

Figure 4.4: Map of the robot control environment with the placement of the four targets. The robot started at target four and the instructions were to reach the targets in order.

Table 4.1: The number of features selected by the GLM feature selection.

Subject\Session	1	2	3
1	1404	1313	1501
2	1265	1132	1214
3	1214	1076	988
4	986	1437	1298

malized to the Montreal Neurological Institute (MNI) reference space and the masks defining the selected voxels were summed over all subjects.

4.3.2 Performance

Table 4.2 shows all subjects' performances in both of the evaluation sessions. The performance was measured by the number of targets reached, the time it took to reach them and in how many movements it was done. All subjects managed to reach at least three of the four targets, and the maximum targets reached was twice that. With 725 images and a TR of 1.62 s, the complete control phase stretched slightly less than 20 min (1175 s).

If we assume that the minimum time between sending one command and the next is 10 TRs (16.2 s, including three TRs for the BOLD signal to reach a detectable

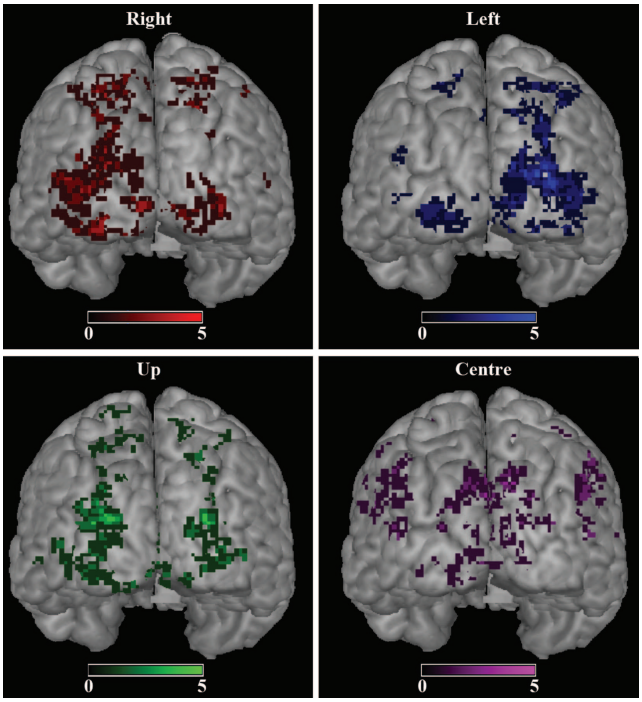

Figure 4.5: Voxels selected in the online GLM analysis, displayed on the Montreal Neurological Institute (MNI) reference brain. Four statistical t-maps, each corresponding to an attention direction, were computed online during data acquisition. From each of these t-maps a mask was created by first locating the 500 highest t-values and then removing any cluster smaller than five voxels. In the online analysis the masks were merged to create the voxel selection to train the SVM on. In this figure, the masks from all subjects were spatially normalized and added, separately for each attention direction. With four subjects and two (performance) sessions each, the sum could take values between 1 and 8. However, since any voxel was selected in at most five sessions, the scale of the overlays has been limited to this value.

level, one TR for the movement and four TRs for the signal to return), the maximum number of commands that can be sent during the experiment is 72. The two best sessions, where five and six targets were reached, did so using 71 and 67 movements respectively. Figure 4.6 visualizes how the robot was maneuvered during the two sessions. Note that only the forward movements result in a new position, i.e. a right turn followed by a left turn cancel out and is not visible in these maps.

Table 4.2: The cumulative time (in seconds) from feedback start to reach the targets. The cumulative number of movements used for reaching the targets is shown within brackets.

		Target					
Subject	**Session**	**1**	**2**	**3**	**4**	**5**	**6**
1	**1**	81 (5)	407 (26)	899 (59)	-	-	-
1	**2**	109 (7)	287 (19)	689 (49)	-	-	-
2	**1**	109 (6)	298 (18)	748 (42)	1042 (58)	-	-
2	**2**	70 (5)	196 (13)	483 (30)	833 (52)	1123 (71)	-
3	**1**	47 (3)	243 (16)	570 (36)	818 (53)	936 (61)	1034 (67)
3	**2**	87 (7)	279 (21)	716 (50)	-	-	-
4	**1**	32 (3)	120 (9)	927 (63)	-	-	-
4	**2**	209 (15)	510 (37)	1011 (61)	-	-	-

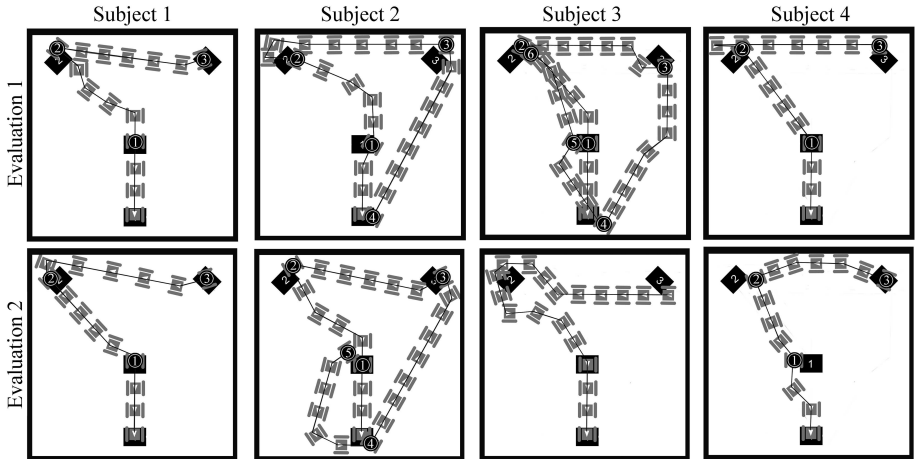

Figure 4.6: The robot's paths during the navigation. The path is only shown until where the last target was reached. Only the forward movements are shown, i.e. right-left-forward looks as simply forward. Each gray "robot symbol" symbolizes the position after a forward movement. The robot's position when hitting a target is indicated with a circle containing the target number.

4.4 Discussion

We have for the first time demonstrated real-time BCI control based on pure covert visuospatial attention, completely independent of eye movements and evoked responses. In a telepresence application, where a robot was moved through a track containing four targets, the user communicated the intended movement by covertly shifting the attention between four different regions in the visual field. Our four subjects were all able

to control the robot and reach at least three of the four targets. All subjects expressed the feeling of having control over the robot, even during the initial practice session. This indicates, in agreement with our findings in Chapters 2 and 3 as well as with other studies [125, 126], that visuospatial attention based BCI control is intuitive with virtually no training required.

Though our study is the first example of an applied BCI based on the visual system that is completely free from evoked responses and eye gaze dependency, the concept of employing the visual system is not new. One example is BCIs based on the steady state visually evoked potential (SSVEP). SSVEP is an evoked response present during a flickering stimulation of the retina, and is detected via an increase of power in the EEG or MEG signal at the frequency of the stimuli. The P300 is another event related potential (ERP) that has been used for BCI. This response occurs approximately 300 ms post-stimulus upon rare events. The matrix speller first described by Farwell et al. [22], is a BCI based on the P300 visual response in EEG signals. Besides being intrinsically dependent on external visual stimulation, there is growing evidence that visual P300 and SSVEP BCI systems are more or less dependent on gaze control [19, 31, 32, 135, 136].

For safety reasons inherent to the high magnetic field, we could not bring an eye tracker into the scanner environment. Thus, we could not get online measures of eye movements. However, it has been shown multiple times that people have no trouble performing covert spatial attention shifts in the absence of any eye movements [49, 51, 106, 123, 124]. More importantly, the brain activity patterns obtained during BCI provide additional evidence (see Chapters 2 and 3) that the subjects controlled the robot via covert shifting of attention, and not with eye movements. It is well known that covert shifting of attention to one side induces elevated activity in the contralateral visual cortex [49, 52, 137]. As can be seen in Figure 4.5, the bulk of activity is contralateral for left and right attentional shifts. If eyes were moved to control the robot, we would expect opposite results, since most of the visual information would shift to the hemifield opposite to the direction of eye movement, causing activity in the visual cortex ipsilateral to that direction. Up and down shifting is associated with inferior and superior visual cortex activation, respectively. Also here the activity patterns are in agreement.

To classify each image volume we trained a support vector machine on the initial localizer data. The application of multivariate classification techniques on fMRI data has been shown effective in multiple studies, e.g. [81, 83, 134]. Since fMRI data samples usually include a very large number of voxels, a feature selection step is most

often included to remove uninformative voxels to avoid overfitting. Our feature selection was based on an online univariate GLM analysis. Four t-maps were computed where in each one, one of the directions was contrasted against the other three. A map of voxels sensitive to at least one of the attention directions was then created by merging the highest values from these maps. It is possible that a multivariate feature selection method could create a map more optimized for the SVM classifier, but our strategy is fast and allowed us to finish the feature selection and training within a single TR. When looking at the overlap of selected voxels across sessions (Figure 4.5), we see that there are regions in expected parts of the cortex that are consistently selected. Around these "hot-spots" there are voxels selected in only a few sessions. There can be several reasons for this spread. First, there is a large variation in peoples' visual field maps [119, 138]. Second, there can be an apparent shift due to misregistrations both during the alignment of the functional data from the two different sessions and during the spatial normalization. Third, there could be small variations in where the subject moves the attention. Subjects reported that they tried different strategies in order to feel confident in directing their attention. These strategies included imagining a beam of light shining from the center onto the target of interest and pretending to expect a symbol to show up at the target. A change of strategy could potentially result in variations of selected voxels.

Ultimately, the goal is to have a BCI system that can function in every-day life for patients, and then MRI is no longer an option. It is only when the setup is successfully moved to a portable system that it can be applied for real. However our results could most likely not be repeated using scalp electrodes, but would demand intracranial recordings. For the control to work it is necessary that the responses to each of the attention directions can be distinguished, both from each other, but also from the visual input from the feedback video. As seen in the activation maps, and as predicted by retinotopic studies, there are multiple cortical regions, corresponding to the multiple visual maps, that get activated during each direction of attention. The brain response to the central input provided by the video camera is closely located and strong compared to the attention modulated effects. Thus, the implicit limitations in terms of resolution and signal strength will probably make EEG ineffective.

Though both EEG and MEG have been used for investigating covert visuospatial attention for BCI control [107, 125, 126], none of these studies perform real-time online decoding and were without visual feedback. It should also be noted that, like fMRI, MEG does not provide a portable technology suitable for the end user. In a recent study [139] Treder et al. used EEG to implement a (ERP dependent) BCI speller

based on both spatial and feature (color) attention, not dependent on eye movements. They evaluated three variants of speller interfaces, two that were sensitive to spatial attention and one that was not. They found the best performance in the version not using spatial attention. For the other two variants, incorporating both spatial and feature attention, the performance dropped substantially when only using the occipital electrodes. This suggests that they did not succeed in detecting the responses to spatial attention.

Intracranial recordings would most likely be suitable for covert attention BCI. In a recent study [46] covert visual attention was studied with ECoG using a classical cueing task. Distinct foci of activity were found indicating that the associated brain signals were readily detectable. Moreover, in Chapter 2 we obtained a performance of 70% with post-hoc offline analysis of ECoG data recorded during a two-direction visual attention task.

Real-time fMRI could be used in further studies aimed at finding the most reliable implant sites that are consistently activated, and to limit the cortical area needed for decoding. As long as their responses can be separated, more directions can be added to achieve a more detailed BCI control. Moreover, the attention target regions have the potential to be moved to positions in the visual field that in turn are mapped to the areas in the cortex where the best signals can be obtained.

In conclusion, we have shown that the visuospatial cortex can be covertly modulated by people to control a robot in real time. The center display did not interfere with the generation of instructions of the robot, suggesting that covert shifting of attention to the periphery can be performed without interfering with processing of information in the center of the field. Conceptually, more than the current three directions can be decoded (diagonal directions or even more), but this needs to be verified.


5

Identification of Cortical Target Regions for Intracranial Brain-Computer Interfacing using Spatially Restricted fMRI Analysis

Abstract

There is a growing interest in Brain-Computer Interfaces (BCI) based on invasive technologies. By implanting electrodes, stronger and more localized signals with higher bandwidth can be recorded. To optimize the signal quality, as well as the surgical procedures, the best electrode placements should be determined before the operation. High-field fMRI measures cortical activity at a high resolution and the metabolic BOLD signals has been shown to correlate well spatially with the cortical potentials recorded with electrodes placed on the brain surface. This makes high-field fMRI exceptionally suited for selecting implant sites. We have in the previous chapters shown that it is possible to decode covertly directed visuospatial attention from individual fMRI images. We here classify attention to four directions using a support vector machine while enforcing several spatial restrictions on the voxels available for the classifier. All the spatial restrictions applied are based on how accessible the brain areas are for implanted surface electrodes. The results show that fMRI signals from only the surface of the brain are sufficient for good BCI performance. Data also show that the topographical pattern is quite variable across subjects. Individual mapping with fMRI is therefore essential for BCI application.

5.1 Introduction

The term Brain-Computer-Interface (BCI) refers to technologies that use measurements of brain activity as control signals to directly communicate with, or physically influence, the outside world. Normally we are dependent on muscular activity for any kind of two-way communication. A direct link between the mind and the physical world promises a better life for paralyzed patients. Until recently BCI research has mainly focussed on scalp recordings using EEG, since implanting electrodes was regarded as prohibitive. With advances in technology, however, interest in intracranial solutions has increased significantly, especially given that fully implantable amplifiers are now available for animals. A result of these new developments is that new options become available for the choice of BCI input signals. Given that the scalp filters electrical signals resulting in loss of higher frequency content and reduced spatial resolution, further research is needed to characterize and validate new target regions and associated brain functions. Studies that decode invasive recordings of brain signals have been done in monkeys [35, 36, 39, 94] for a longer time, but since the late nineties there have also been studies on humans [38, 40, 43, 140–142], with positive results. Electrodes can be implanted to record from the cortical surface, electrocorticography (ECoG), or inside the cortex using intracortical electrode arrays. In order to achieve a good result, it is important to find techniques to do as much preparation and optimization as possible before the surgery, including pre-localizing the best electrode placements. This will allow for minimally invasive surgery and it will minimize the risk of the patient not being able to produce a reliable control signal. Since fMRI is the only functional imaging modality that non-invasively gives access to the whole brain at a high resolution, it is well suited to be used for this purpose. It has been shown that the BOLD changes measured with fMRI correlate well spatially with direct electrophysiological measurements using ECoG [43, 89–91]. This seems to be particulary true for the power in the gamma band (65-95 Hz) [89, 91].

It has been well established that covert visuospatial attention, while maintaining central fixation, induces activity mainly in the same parts of the cortex as a visual stimulus at the attended region [48–55]. The visual cortex contains multiple maps, representing the visual field, located in the occipital cortex. Therefore, when a particular point of the retina is stimulated or attended to, there is a response in multiple regions of the cortex.

In Chapter 3 we showed that real-time fMRI can be used to decode visuospatial attention to four directions. Since we allowed all regions in the occipital cortex to be used in that analysis, the results were potentially partly based on activity in deeper

brain regions that are not accessible with surface electrodes. In the current chapter we investigate whether BCI with covert visuospatial attention is feasible with recordings restricted to the accessible regions. The fMRI data from Chapter 3 were reanalyzed to study the effects of excluding voxels located in certain brain areas on BCI performance. These restrictions are discussed in the following paragraphs.

Firstly, owing to the risk of damaging vessels it is usually hard, or even impossible, to place ECoG electrodes inside a sulcus. The electrode grids are therefore placed on the "convex hull" defined by the gyri. Accordingly, a restriction we applied in our analysis was to only include voxels located on the occipital surfaces.

Secondly, the interhemispheric medial occipital surfaces are more difficult to implant on compared to the lateral surfaces because of larger veins and the limited space between the cortex and the falx. We therefore applied the restriction of excluding the interhemispheric surfaces.

Thirdly, we restricted the volume of interest to a single hemisphere. The motivation is that it would be beneficial if one could limit the surgery to a single cranial entry point, but this would mean that one only has access to one of the occipital lobes. Since the left and right hemifields are mainly processed in the contralateral hemisphere, one might expect that using a single entry point will have a large effect on the BCI performance.

The final restriction we made was to exclude BOLD activation clusters smaller than a predefined size. This was done to make sure that the classification was not based on many single voxels, or small voxel clusters, spread over the cortex. In this case, it would be difficult to sufficiently cover the areas of interest with ECoG electrodes.

By analyzing the fMRI data with the different spatial restrictions based on issues and limitations of ECoG, we could make the results more relevant for the later modality. By comparing the classification performances we could estimate the implications of the restrictions on the performance and make the case for visuospatial attention being an effective target in ECoG based BCI stronger.

5.2 Methods and Materials

In this chapter we reanalyze the data presented in Chapter 3. A brief description of the experimental setup and analysis is given here. More details can be found in Chapter 3.

5.2.1 Subjects

Nine healthy volunteers (age 23-28, right-handed, 5 male) with normal or corrected-to-normal vision and naive to the task participated in the study. The study was approved by the ethics committee of the University Medical Center Utrecht in accordance with the declaration of Helsinki (2008) and all subjects had given their written informed consent.

5.2.2 Data

The subjects were scanned with a 7T Philips Achieva system (Philips Healthcare, Cleveland, OH, USA) with a 16-channel SENSE head coil (Nova Medical, MA, USA). The functional data were recorded using an EPI sequence (TR/TE=1620/25ms; FA=90; SENSE factor=2; 35 coronal slices, acquisition matrix 96x96, slice thickness 2mm with no gap, 1.848 mm in-slice resolution). The FOV was selected such that it covered the occipital lobe and the most posterior part of the parietal lobe. Each experiment consisted of 652 volumes from which the first 292 volumes were used as training data and the remaining part served as testing data.
A high-resolution image was acquired for the anatomy using a T_1 3D TFE sequence (TR/TE=6/2ms; FA=7; FOV=220x180x200mm; 0.55x0.55x0.5mm reconstructed resolution).

5.2.3 Task

During the experiment the subjects were presented with an image as seen in Figure 3.1b. The image was projected via a projector into the bore of the scanner and onto a small projection screen. The image contained a central cue surrounded by four yellow arrows. The central cue was both the point on which the subjects kept their gazed fixed at all times and where instructions of where to move the attention were communicated. The cue either showed a circle or an arrow pointing towards one of the yellow target areas. Subjects were instructed to covertly move the attention to the target area located in the direction of the instruction arrow or, in the case of a circle, to relax and keep attending to the center.

The setup of the image projection limited the size of the image and it was more restricted vertically than horizontally. The angular distance from the central point was 6° to the up and the down regions and 11° to the left and right regions.

5.2.4 Data Preprocessing

All volumes in the fMRI run were aligned to the first one. The rigid registration was performed by minimization of the sum of squared differences between gray-value intensities. A stochastic gradient descent method [111] was employed for optimization, using 50 iterations. Prior to the registration, the images were blurred with a Gaussian filter (σ = 1 voxel). Linear interpolation was used during the optimization steps while the final image was generated using cubic B-spline interpolation. The algorithm was implemented in C++, and called from Matlab (Mathworks, Natick, MA).

We also applied detrending to the training data before the classifier training, and to each volume before it was classified (using only the preceding images). We used an implementation of the algorithm described in [114] with regularization parameter $\lambda = 200$. Finally, each voxel's signal was normalized to have zero mean and unit standard deviation, as estimated from the training data.

5.2.5 Segmentation

Each restriction of voxels included in the analysis was represented by a binary mask (except the cluster-size restriction that took place during the feature selection step).

The first image of an fMRI run was manually segmented. The whole imaged part of the brain was extracted and this mask was then separated into the left and right hemispheres. The reason for using the first image is that it served as template for motion correction.

Surface masks for the single hemispheres were created by first subtracting a morphologically eroded version of the solid mask, and from this removing the frontal coronal "wall" that is not representing a true surface, but is rather an artifact from the fact that we do not scan the whole brain. To make sure the medial surfaces were included, the full surface was created by taking the union of the single hemisphere surfaces. We also made surface masks excluding the medial surfaces. The medial sides of the occipital lobes are harder to reach with surface electrodes and we wanted to see if good decoding was still possible when they were excluded.

The structuring element used for the erosion was a sphere with a radius of 3 voxels. With a resolution of 2 mm this resulted in a mask selecting the 6 mm closest to the surface. Note that here the surface does not mean the most superficial cortical layer, but rather the convex hull of the brain determined by its gyri (Figure 5.1). This method agrees with our assumption that implanting electrodes in sulci is not an option.

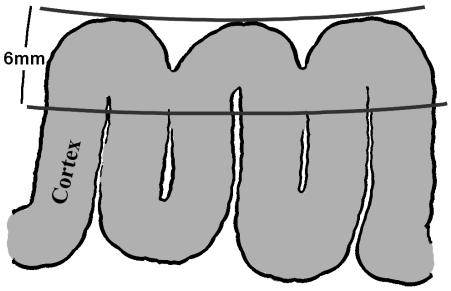

Figure 5.1: Schematic of a surface mask restricting the analysis to the voxels located within 6 mm of the surface.

5.2.6 Classifier

As classifier we applied the C-SVM classifier with a linear kernel and regularization parameter $C = 1$ from the PyMVPA python software package [143]. LaConte et al. [134] have shown that the C-SVM is quite insensitive to the value of C, and that this value works well. Since an SVM can only model a binary decision problem, and we have a four class problem, our classifier actually consisted of six SVMs where each separated two of the classes. A tested image volume was assigned to the class with the majority vote among these six binary classifiers.

5.2.7 Voxel selection

With the low ratio of data samples to number of voxels, it is necessary to apply some method to reduce the number of voxels included in the classification model in order to avoid overfitting. A common strategy used in machine learning when wanting to reduce the dimensionality of the problem by removing features is to first rank them in order of importance and then remove the ones with the lowest ranking. We applied this strategy to our data to locate and select the most relevant voxels. The ranking was done based on sensitivity maps computed using two different methods: a univariate general linear model (GLM) based method and a multivariate method using the same SVM classifier as during the classification. The first method is very fast and can therefore perform the selection in parallel with the data collection. For this reason we have applied it in our previous studies (Chapters 2 to 4). In contrast with the previous chapters, we here use a multivariate classifier and perform the analysis offline. The

second method, a multivariate sensitivity analysis, was thus included to test whether taking the SVM classifier into account could provide a voxel selection that improved the classification results. A number of the most informative voxels were selected as features to train the classifier on. The different spatial restrictions we applied to the data resulted in a different volume of cortex and thus the number of voxels that were selected was adjusted accordingly (see Section 5.2.7.3).

5.2.7.1 GLM sensitivity maps

This voxel selection is based on t-maps from a univariate GLM analysis. Four statistical t-maps were computed using five regressors; right, left, up and down attention plus a linear drift term. The t-values were computed using the contrasts 'one minus the others'. That is, for right attention the contrast was 'right - $\frac{1}{3}$[left+up+down]' etc. The voxels were ranked based on these t-values and a selection was made by merging the voxels with the highest absolute values from each of the four t-maps, see Section 5.2.7.3.

5.2.7.2 Multivariate sensitivity maps

Since we use a linear SVM classifier, a simple measure of a voxel's relevance is the magnitude of its weight in the trained model [144]. Since the classifier consisted of six binary SVMs, there were six sensitivity maps. The final map was created by taking the maximum sensitivity across the maps for each voxel. This map indicated whether or not the voxel was important for separating any of the pairs of classes.

A possibly improved selection could be reached by using a recursive feature elimination scheme, but this would be at the expense of a much longer computation time.

5.2.7.3 Number of voxels

From the multivariate sensitivity maps a fixed number of voxels were selected. However, when using the GLM method, to ensure that we had a fair representation of voxels sensitive to each attention direction, we first selected a quarter of this number from each of the four t-maps and then merged these voxels into a single selection. Any overlap meant that the merged number of voxels was lower than the fixed number.

The number of voxels selected from the solid two-hemispheres mask was 1500 (corresponding to 10 cm^3). For the solid left and right hemisphere masks the number of voxels selected was 750. The volume of the surface mask was roughly a third (average 37%) of that of the full mask. Accordingly, the number of voxels selected

was 500 for the two-hemispheres surface and 250 for the single hemisphere surfaces. We did not make a difference in the number of voxels used for the surface masks when the medial surfaces were included or excluded.

From all these selections we removed clusters smaller than 5 voxels. Without scattered small clusters, the selection better represents a possible ECoG coverage.

5.2.8 Training and Classification

The classifier was trained on the first 292 volumes that consisted of nine trials for each direction in randomized order and separated by a rest condition. Each trial, as well as the rest condition, was 6.48s (4 TRs) long. The test data consisted of ten trials of 9.72s (6 TRs) for each of the four directions, and were separated by a rest condition of 4.86s (3 TRs). Each trial was classified using only the fifth volume, giving enough time for the hemodynamic response to reach a detectable level (see Chapter 2 and 3).

5.3 Results

5.3.1 Segmentation

Figure 5.2 shows an example of the two-hemispheres surface mask. The figure displays both a coronal and a sagittal slice of EPI data with the surface mask overlaid in gray and white. The gray part represents the voxels not included in the mask when the medial surfaces are left out.

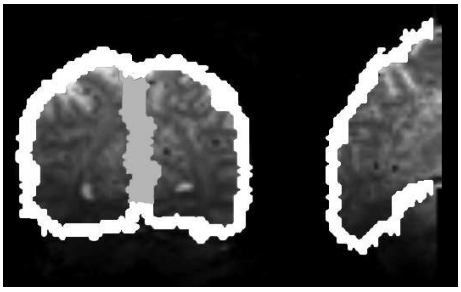

Figure 5.2: A coronal (left) and sagittal (right) slice of EPI data with the two hemispheres surface mask. The gray part of the mask represents the medial surfaces that may be included or excluded in the mask.

5.3.2 Number of voxels selected from the GLM sensitivity maps

The method used in the GLM based voxel selection allowed a voxel to be selected from more than one of the four t-maps. When the selections were merged to create the final set, any overlap meant that there were fewer voxels than the number defined for the particular mask type. Table 5.1 shows the final number of voxels included in the GLM based selections. Only the numbers corresponding to the surfaces excluding the medial sides are included in the table.

Table 5.1: The number of voxels selected by the GLM based voxel selection. The maximum number, same as the number selected from the multivariate sensitivity map, is shown within brackets. (*full* - both left and right hemispheres, *LH/RH* - left/right hemisphere, *surf* - surface excluding medial parts)

Subject	**full** (1500)	**full$_{surf}$** (500)	**LH** (750)	**LH$_{surf}$** (250)	**RH** (750)	**RH$_{surf}$** (250)
1	1161	359	599	184	564	166
2	1245	424	603	199	665	223
3	1138	388	579	189	548	172
4	1196	397	619	193	570	192
5	1135	383	513	168	588	199
6	1195	420	626	213	565	200
7	1182	405	623	202	571	197
8	1134	389	557	179	605	214
9	1145	411	567	209	558	203
Avg	1170	397	587	193	582	196

5.3.3 Number of voxels after removing clusters< 5

We made additional feature selection masks by removing any cluster smaller than five voxels from the original selections. Table 5.2 shows the number of voxels that were left after this exclusion criteria was applied. Note that for the GLM based selections the original numbers are the ones seen in Table 5.1, while for the multivariate sensitivity method, the starting point was the full number defined for the mask (within brackets in Table 5.1). As in Table 5.1, we only included the numbers corresponding to the surfaces excluding the medial parts.

Figure 5.3 shows for all subjects the voxels selected by the multivariate sensitivity analysis when the choice was restricted to the non-medial surface of both hemispheres and the small clusters had been removed.

Table 5.2: The number of voxels selected from the GLM and the multivariate (SVM) sensitivity maps, after removing clusters smaller than five voxels.

	GLM						SVM					
Subject	**full**	**full$_{surf}$**	**LH**	**LH$_{surf}$**	**RH**	**RH$_{surf}$**	**full**	**full$_{surf}$**	**LH**	**LH$_{surf}$**	**RH**	**RH$_{surf}$**
1	867	242	414	115	395	111	1122	357	527	175	530	154
2	906	251	418	118	405	88	827	279	369	144	367	111
3	1013	302	429	134	460	119	1345	397	610	190	649	179
4	929	280	474	135	401	119	1167	356	560	191	527	161
5	1008	281	429	138	492	136	1319	372	597	179	584	166
6	959	281	399	115	465	131	925	251	387	102	409	105
7	896	284	450	118	471	146	1020	317	412	148	480	129
8	982	304	480	141	485	167	1214	327	561	160	597	170
9	1047	356	508	178	518	179	1387	443	657	205	681	225
Avg	956	287	445	132	455	133	1147	344	520	166	536	156

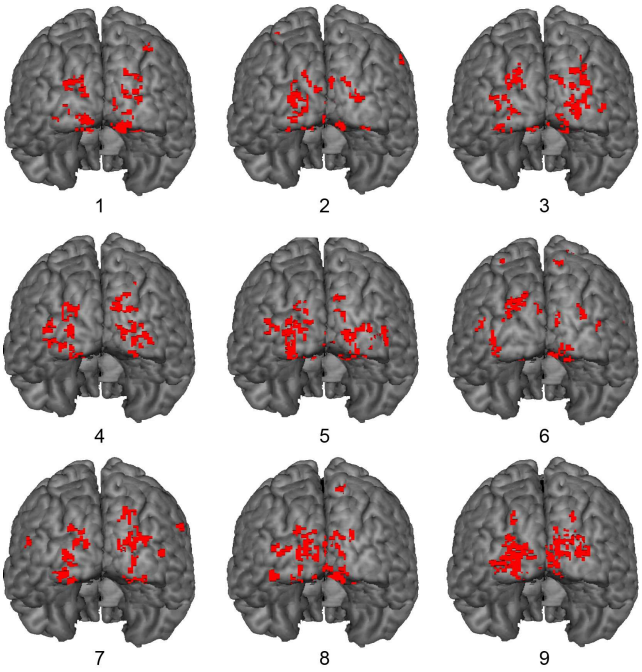

Figure 5.3: Voxels selected by the SVM sensitivity analysis when restricted to the surface of both hemispheres (excluding the medial sides) and after removal of clusters smaller than 5 voxels. Each subject's selection has been normalized to the Montreal Neurological Institute (MNI) reference space.

5.3.4 Classification performance

Each trial, consisting of six image volumes, was classified based only on the fifth image volume. Because of the delay in the hemodynamic response it takes a few seconds before the neural activity can be detected in the BOLD signals. In keeping with the real-time data processing during the experiments the instruction was updated on average 1.1 seconds into the trial (see Chapter 3). This means that we classified the BOLD effect that was present seven seconds after the attention shift, which has been shown to be enough for an effective decoding (see Chapters 2 and 3).

5.3.4.1 Voxel selection from GLM sensitivity maps

Tables 5.3-5.5 show the classification performance when the SVM was trained on voxels selected based on the GLM sensitivity maps. The tables are divided in three vertical and two horizontal sections. The vertical sections represent the solid volume and the surfaces including and excluding the medial sides, while the horizontal sections divide the results into when allowing all cluster sizes and when only using clusters larger than five voxels.

The results when using voxels from both hemispheres are shown in Table 5.3. On average 93.1% of the trials were classified correctly when using the solid hemispheres with restricted cluster sizes. Only using the surface voxels lowered the number to 89.7 or 87.5%, depending on whether or not the medial surface was included.

The performances when restricting the available voxels to only the left or the right hemisphere are shown in Table 5.4 and Table 5.5, respectively. As expected, these restrictions lowered the performance for the ipsilateral attention. For instance, using only the solid right hemisphere, the performance for right attention drops 17 percentage points compared to when using both. The drop was slightly smaller, around 10 percentage points for left attention in the corresponding scenario. Also the classification of up and down attention was affected by removing a hemisphere, but the effect is only about 5 percentage points. Even with the hardest restriction, the surface of the single left or right hemisphere excluding the medial part and small clusters removed, the average performance over all four directions was 82.2% and 76.7%, respectively. The overall effect on performance of excluding small clusters was small. When using the solid volumes there was a small positive effect, while the performance dropped 2-4 percentage points when the analysis was restricted to the surfaces. The reason for this difference in the effect is probably that when only using the surfaces you start with such a small volume that each included voxel has a larger impact on the classification.

Table 5.3: Percentage of the ten trials for each attention direction that were classified correctly when the voxels were selected from both hemispheres based on the GLM sensitivity maps. Each trial was classified based only on the fifth image volume. (R,L,U,D = right,left,up,down attention)

		all clusters					clusters$>$ 5				
	Subject	**R**	**L**	**U**	**D**	**Avg**	**R**	**L**	**U**	**D**	**Avg**
Solid	1	80	90	90	80	85.0	80	90	90	90	87.5
	2	100	90	90	90	92.5	100	90	90	90	92.5
	3	100	100	100	100	100.0	100	100	100	100	100.0
	4	90	100	100	100	97.5	90	100	100	100	97.5
	5	90	100	100	100	97.5	90	90	100	100	95.0
	6	90	100	100	80	92.5	90	100	100	80	92.5
	7	100	90	90	90	92.5	100	100	90	90	95.0
	8	90	100	100	90	95.0	90	100	100	90	95.0
	9	70	60	100	100	82.5	70	60	100	100	82.5
	Avg	90.0	92.2	96.7	92.2	92.8	90.0	92.2	96.7	93.3	93.1
Surface incl med	1	60	90	90	80	80.0	70	90	80	90	82.5
	2	90	90	100	90	92.5	100	90	90	100	95.0
	3	100	100	100	100	100.0	100	100	100	100	100.0
	4	90	100	100	100	97.5	90	100	100	100	97.5
	5	90	90	100	90	92.5	90	80	100	90	90.0
	6	70	100	90	70	82.5	80	100	90	70	85.0
	7	90	80	90	90	87.5	90	90	100	80	90.0
	8	80	90	90	100	90.0	90	80	90	100	90.0
	9	70	60	100	100	82.5	60	60	90	100	77.5
	Avg	82.2	88.9	95.6	91.1	89.4	85.6	87.8	93.3	92.2	89.7
Surface excl med	1	60	90	90	90	82.5	60	90	80	90	80.0
	2	90	90	100	100	95.0	100	90	80	80	87.5
	3	100	100	100	100	100.0	100	100	100	100	100.0
	4	90	100	100	100	97.5	90	100	100	100	97.5
	5	90	100	100	90	95.0	90	80	100	90	90.0
	6	70	100	100	70	85.0	90	100	100	70	90.0
	7	90	90	90	80	87.5	90	90	90	60	82.5
	8	80	80	90	100	87.5	60	90	90	100	85.0
	9	70	50	100	100	80.0	60	60	90	90	75.0
	Avg	82.2	88.9	96.7	92.2	90.0	82.2	88.9	92.2	86.7	87.5

Table 5.4: Percentage of the ten trials for each attention direction that were classified correctly when the voxels were selected from the left hemispheres based on the GLM sensitivity maps. Each trial was classified based only on the fifth image volume. (R,L,U,D = right,left,up,down attention)

		all clusters					clusters>5				
	Subject	**R**	**L**	**U**	**D**	**Avg**	**R**	**L**	**U**	**D**	**Avg**
Solid	1	80	80	80	70	77.5	80	80	90	70	80.0
	2	100	60	80	90	82.5	100	40	90	100	82.5
	3	90	90	100	100	95.0	90	90	100	100	95.0
	4	70	90	90	100	87.5	80	90	90	100	90.0
	5	100	80	100	90	92.5	100	80	100	90	92.5
	6	90	80	90	90	87.5	100	90	100	70	90.0
	7	100	100	80	70	87.5	90	90	80	80	85.0
	8	80	90	100	90	90.0	70	80	100	90	85.0
	9	70	70	100	90	82.5	80	70	100	90	85.0
	Avg	86.7	82.2	91.1	87.8	86.9	87.8	78.9	94.4	87.8	87.2
Surface incl med	1	80	80	80	80	80.0	70	90	70	80	77.5
	2	90	30	80	90	72.5	100	40	80	100	80.0
	3	90	80	100	100	92.5	100	80	100	100	95.0
	4	80	100	90	100	92.5	80	90	60	90	80.0
	5	90	80	100	90	90.0	80	70	100	90	85.0
	6	100	90	100	90	95.0	100	80	100	90	92.5
	7	100	60	90	70	80.0	90	50	80	70	72.5
	8	90	80	100	90	90.0	80	80	90	90	85.0
	9	80	50	100	90	80.0	70	40	100	80	72.5
	Avg	88.9	72.2	93.3	88.9	85.8	85.6	68.9	86.7	87.8	82.2
Surface excl med	1	80	80	80	80	80.0	80	80	70	80	77.5
	2	80	40	80	100	75.0	90	60	90	100	85.0
	3	100	80	100	90	92.5	90	80	100	100	92.5
	4	80	100	90	100	92.5	80	90	60	90	80.0
	5	90	80	100	90	90.0	80	70	100	90	85.0
	6	100	80	100	70	87.5	100	70	100	70	85.0
	7	100	50	90	70	77.5	90	50	80	70	72.5
	8	80	70	100	90	85.0	80	80	90	90	85.0
	9	80	60	100	90	82.5	80	40	100	90	77.5
	Avg	87.8	71.1	93.3	86.7	84.7	85.6	68.9	87.8	86.7	82.2

Table 5.5: Percentage of the ten trials for each attention direction that were classified correctly when the voxels were selected from the right hemispheres based on the GLM sensitivity maps. Each trial was classified based only on the fifth image volume. (R,L,U,D = right,left,up,down attention)

		all clusters					clusters$>$ 5				
	Subject	**R**	**L**	**U**	**D**	**Avg**	**R**	**L**	**U**	**D**	**Avg**
Solid	1	40	90	90	50	67.5	50	90	100	90	82.5
	2	60	90	80	70	75.0	60	90	80	80	77.5
	3	90	100	100	100	97.5	90	100	100	100	97.5
	4	100	90	100	100	97.5	100	90	100	100	97.5
	5	100	90	80	90	90.0	80	90	80	90	85.0
	6	50	90	100	80	80.0	50	90	100	90	82.5
	7	60	100	80	80	80.0	70	100	80	60	77.5
	8	90	60	90	100	85.0	90	60	90	100	85.0
	9	70	90	90	100	87.5	70	90	90	100	87.5
	Avg	73.3	88.9	90.0	85.6	84.4	73.3	88.9	91.1	90.0	85.8
Surface incl med	1	60	60	90	50	65.0	40	80	90	60	67.5
	2	70	80	80	50	70.0	80	80	20	50	57.5
	3	90	100	100	100	97.5	80	100	90	100	92.5
	4	100	90	100	100	97.5	100	80	100	100	95.0
	5	100	80	80	90	87.5	80	90	80	90	85.0
	6	30	90	90	80	72.5	40	80	100	80	75.0
	7	60	90	70	60	70.0	60	90	70	60	70.0
	8	50	50	70	100	67.5	70	40	70	90	67.5
	9	80	90	80	90	85.0	60	90	80	80	77.5
	Avg	71.1	81.1	84.4	80.0	79.2	67.8	81.1	77.8	78.9	76.4
Surface excl med	1	60	60	90	50	65.0	40	80	90	60	67.5
	2	70	90	80	40	70.0	70	90	40	50	62.5
	3	90	100	100	100	97.5	90	100	90	90	92.5
	4	100	90	100	100	97.5	100	80	100	100	95.0
	5	100	90	90	90	92.5	80	90	70	90	82.5
	6	60	90	90	80	80.0	60	90	70	70	72.5
	7	60	90	60	80	72.5	60	90	70	50	67.5
	8	60	40	70	90	65.0	80	40	70	100	72.5
	9	80	90	80	90	85.0	60	90	80	80	77.5
	Avg	75.6	82.2	84.4	80.0	80.6	71.1	83.3	75.6	76.7	76.7

5.3.4.2 Voxel selection from multivariate sensitivity maps

Tables 5.6-5.8 show the classification performance when the voxels were selected based on the multivariate sensitivity maps. Trained on these selections the classifier performed better than when the selections were based on the univariate method. The correctly classified trials were now 95.3% (Table 5.6) when using both solid hemispheres but restricting the cluster sizes. Restricting the selection to the surface voxels reduced the performance by five percentage points, i.e. a similar effect as for the GLM based method. The drop in performance when excluding one of the hemispheres was not as large. With only the single left or right hemisphere, about 88% of the trials were still classified correctly. With the overall better classification using this method the average performances were 85.0% and 78.9% when only allowing voxels from the non-medial surface of the single left or right hemisphere and with small clusters removed.

Table 5.6: Percentage of the ten trials for each attention direction that were classified correctly when the voxels were selected from both hemispheres based on the multivariate sensitivity maps. Each trial was classified based only on the fifth image volume. (R,L,U,D = right,left,up,down attention)

		all clusters					clusters> 5				
	Subject	**R**	**L**	**U**	**D**	**Avg**	**R**	**L**	**U**	**D**	**Avg**
Solid	1	100	100	80	90	92.5	100	100	90	90	95.0
	2	100	90	100	80	92.5	100	100	100	90	97.5.0
	3	100	100	100	100	100.0	100	100	100	100	100.0
	4	90	100	100	100	97.5	90	100	100	100	97.5.0
	5	100	90	100	100	97.5	100	90	100	100	97.5
	6	80	90	100	90	90.0	90	90	100	90	92.5.0
	7	100	100	80	70	87.5	100	100	90	80	92.5
	8	100	100	100	100	100.0	100	100	100	100	100.0
	9	70	90	100	100	90.0	70	70	100	100	85.0
	Avg	93.3	95.6	95.6	92.2	94.2	94.4	94.4	97.8	94.4	95.3
Surface incl med	1	80	90	90	90	87.5	80	80	90	90	85.0
	2	90	90	100	80	90.0	100	90	70	80	85.0
	3	100	100	100	100	100.0	100	100	100	100	100.0
	4	90	100	100	100	97.5	100	90	100	100	97.5
	5	100	90	100	90	95.0	100	90	100	90	95.0
	6	80	100	100	80	90.0	80	100	80	90	87.5
	7	90	80	80	100	87.5	90	90	100	100	95.0
	8	90	90	80	100	90.0	90	90	80	100	90.0
	9	70	80	100	90	85.0	70	70	90	90	80.0
	Avg	87.8	91.1	94.4	92.2	91.4	90.0	88.9	90.0	93.3	90.6
Surface excl med	1	70	90	90	90	85.0	60	80	90	90	80.0
	2	90	90	90	90	90.0	100	90	80	90	90.0
	3	100	100	100	100	100.0	100	100	100	100	100.0
	4	90	100	100	100	97.5	100	90	100	100	97.5
	5	100	90	100	90	95.0	100	90	100	90	95.0
	6	100	100	100	80	95.0	90	100	90	90	92.5
	7	90	90	80	100	90.0	90	90	90	90	90.0
	8	80	70	80	100	82.5	70	70	80	100	80.0
	9	70	70	100	90	82.5	70	70	90	100	82.5
	Avg	87.8	88.9	93.3	93.3	90.8	86.7	86.7	91.1	94.4	89.7

Table 5.7: Percentage of the ten trials for each attention direction that were classified correctly when the voxels were selected from the left hemispheres based on the multivariate sensitivity maps. Each trial was classified based only on the fifth image volume. (R,L,U,D = right,left,up,down attention)

		all clusters					**clusters$>$ 5**				
	Subject	**R**	**L**	**U**	**D**	**Avg**	**R**	**L**	**U**	**D**	**Avg**
Solid	1	90	80	80	80	82.5	70	80	80	90	80.0
	2	100	70	100	70	85.0	100	80	100	70	87.5
	3	90	90	100	100	95.0	90	90	100	100	95.0
	4	70	90	90	100	87.5	80	100	90	100	92.5
	5	100	90	100	100	97.5	100	90	100	100	97.5
	6	100	70	100	90	90.0	100	60	100	100	90.0
	7	100	70	80	80	82.5	90	60	80	90	80.0
	8	90	80	100	100	92.5	90	70	100	100	90.0
	9	70	70	100	90	82.5	80	70	100	90	85.0
	Avg	90.0	78.9	94.4	90.0	88.3	88.9	77.8	94.4	93.3	88.6
Surface incl med	1	90	80	90	90	87.5	60	80	90	80	77.5
	2	100	40	90	80	77.5	90	50	100	100	85.0
	3	90	90	100	100	95.0	90	90	100	100	95.0
	4	90	100	90	100	95.0	80	100	90	100	92.5
	5	100	80	100	90	92.5	100	70	100	90	90.0
	6	100	80	90	80	87.5	100	90	70	90	87.5
	7	100	60	70	100	82.5	90	50	90	90	80.0
	8	80	70	100	100	87.5	80	50	80	90	75.0
	9	90	60	100	90	85.0	90	40	100	90	80.0
	Avg	93.3	73.3	92.2	92.2	87.8	86.7	68.9	91.1	92.2	84.7
Surface excl med	1	90	70	80	80	80.0	70	70	90	90	80.0
	2	100	60	100	90	87.5	90	70	100	90	87.5
	3	90	90	100	100	95.0	90	90	100	90	92.5
	4	80	90	80	100	87.5	80	100	90	100	92.5
	5	100	80	100	100	95.0	100	70	100	90	90.0
	6	90	80	100	80	87.5	100	80	100	80	90.0
	7	90	70	80	100	85.0	90	50	80	90	77.5
	8	90	60	90	90	82.5	70	60	80	90	75.0
	9	90	70	90	90	85.0	90	40	100	90	80.0
	Avg	91.1	74.4	91.1	92.2	87.2	86.7	70.0	93.3	90.0	85.0

Table 5.8: Percentage of the ten trials for each attention direction that were classified correctly when the voxels were selected from the right hemispheres based on the multivariate sensitivity maps. Each trial was classified based only on the fifth image volume. (R,L,U,D = right,left,up,down attention)

		all clusters					clusters$>$ 5				
	Subject	**R**	**L**	**U**	**D**	**Avg**	**R**	**L**	**U**	**D**	**Avg**
Solid	1	40	100	90	80	77.5	60	90	100	80	82.5
	2	70	90	80	90	82.5	70	80	80	80	77.5
	3	90	100	100	100	97.5	90	100	100	100	97.5
	4	100	90	100	100	97.5	100	90	100	100	97.5
	5	90	90	70	90	85.0	100	90	90	90	92.5
	6	60	80	100	80	80.0	50	80	100	90	80.0
	7	50	100	60	60	67.5	80	100	80	60	80.0
	8	100	90	90	100	95.0	90	100	90	100	95.0
	9	80	90	90	100	90.0	70	90	90	100	87.5
	Avg	75.6	92.2	86.7	88.9	88.8	78.9	91.1	92.2	88.9	87.8
Surface incl med	1	70	90	90	80	82.5	40	70	100	80	72.5
	2	40	90	80	70	70.0	60	90	70	50	67.5
	3	90	100	100	100	97.5	90	100	100	100	97.5
	4	100	90	100	100	97.5	90	90	100	100	95.0
	5	80	90	90	90	87.5	100	90	90	90	92.5
	6	70	80	90	70	77.5	40	90	70	80	70.0
	7	30	90	70	60	62.5	30	90	60	60	60
	8	100	50	80	100	82.5	80	70	80	80	77.5
	9	80	90	90	100	90.0	80	90	90	100	90
	Avg	73.3	85.6	87.8	85.6	83.1	67.8	86.7	84.4	82.2	80.3
Surface excl med	1	60	90	100	80	82.5	40	70	100	80	72.5
	2	50	100	90	70	77.5	70	100	30	40	60.0
	3	90	100	100	100	97.5	90	100	100	100	97.5
	4	100	80	100	100	95.0	100	80	100	100	95.0
	5	90	90	90	90	90.0	90	80	90	90	87.5
	6	60	90	80	70	75.0	60	80	40	80	65.0
	7	40	90	80	70	70.0	30	90	80	70	67.5
	8	70	70	80	90	77.5	80	70	80	70	75.0
	9	80	90	90	100	90.0	80	90	90	100	90.0
	Avg	71.1	88.9	90.0	85.6	83.9	71.1	84.4	78.9	81.1	78.9

5.4 Discussion

By performing the same classification analysis with different criteria on voxel inclusion, we have been able to estimate the implications of the anatomical limitations that have to be considered when planning to implant electrodes. In spite of significant anatomical constraints we could decode the directed attention with a high accuracy.

When allowing voxels anywhere in the scanned part of the brain to be included, 95% of the trials were classified correctly. When limiting the analysis to the voxels located within 6 mm of the surface, the performance was still as high as 91%. These voxels represent the part of cortex that can be monitored by surface electrodes, and the maintained high performance suggest that ECoG based BCI based on visuospatial attention is possible. It is harder to position electrodes on the medial surfaces, between the cerebral hemispheres, compared to the lateral surfaces. We therefore wanted to test if it is crucial to have electrodes placed on these parts of the cortex. Despite the fact that a large part of the visual cortex, especially V1, is located here our analysis shows that these difficult areas can be excluded without sacrificing classification quality. The difference in performance between including and excluding the medial surface was only about one percentage point.

As expected, when using only one of the hemispheres, classifying attention to the ipsilateral side became much harder. The other three directions were affected much less and could still be recognized in 80-90% of the trials.

The exclusion of clusters consisting of less than five voxels did not have a big impact on the performance. Still, as Table 5.2 shows, an average of 28% of the selected voxels were removed in this step.

There was a substantial improvement in the classification results when the voxel selection was based on the multivariate sensitivity maps. When the SVM was trained on the voxels selected by the GLM based method, the performance was lower overall. Even though the classification performance never differed by more than four percentage points when applying the two selection methods, these results suggest that a lot can be gained by using a more sophisticated voxel selection method. This will likely entail that the analysis can no longer keep up with the scanning and that the acquisition of training and testing data need to be split into two separate fMRI runs. Computation time is then less of an issue and iterative selection schemes such as Recursive Feature Elimination [85] could be applied.

Both the sensitivity analysis and the cluster size restriction in this study were done using simple methods that were computationally effective. Further studies need to be done to optimize these steps in terms of giving the best performance under the restric-

tions motivated by e.g. ECoG. Ideally, the cluster analysis will be part of an iterative feature selection scheme and will not be done in a separate step. Despite the relative simplicity of the methods used here, we reached an average correct classification rate of 85% when only allowing clusters larger than five voxels from the non-medial surface of the left hemisphere (79% for the right hemisphere). Table 5.2 shows that with these selections the average number of voxels used for classification was only 166 (156 for the right hemisphere), corresponding to a tissue volume of 1.1 cm^3. When including the surface of both hemispheres the volume of selected voxels was 2.3 cm^3, still a very small volume, and the average performance was as high as 90%.

The fact that we can decode four "commands" within the small volume of cortex, located in occipital areas accessible by surface electrodes, clearly shows the potential of visuospatial attention as a method for BCI control. There are two main benefits to the small area needed to decode the attention shifts. Firstly, there is less area that needs to be covered by implanted electrodes. Secondly, there is probably less risk of false positives since attention to (or stimuli in) regions close to the attention target areas will not affect the signals in the covered cortex.

A closer look revealed that when restricting the voxel selection to the left hemisphere surface the (on average 166) voxels were distributed over an average of eight clusters, or two clusters per attention direction. Though further studies are needed to confirm this, a more sophisticated voxel selection and cluster analysis could probably limit the selection to a single cluster, and in the extension electrode, per direction.

Figure 5.3 shows that while there are clear similarities in activation patterns over subjects on a large scale, there are also large variations in the precise locations of the clusters. This distribution is expected since it is known that there is no anatomical consistency in the topography of the visual regions [119, 138]. This means that a standard atlas can not provide detailed information regarding the placement of electrodes. It is therefore important to have access to the individual high quality mappings of the visual cortex that can be produced using high field fMRI.

Our visuospatial attention task is attractive for BCI control since it does not depend on any evoked response or eye movements. Owing to the high magnetic field it was not safe to use an eye-tracker to record online measures of eye movements. In our previous presentation of the current data (Chapter 3) we showed how the activation patterns themselves can motivate that the gaze was fixed at the center. In short, moving the gaze towards the target will create an activation pattern mirrored to covert attention shifts. It should also be mentioned that it has been shown in several other studies on covert visual attention that shifts can easily be performed in the absence of any eye

movements [49, 51, 106, 123, 124].

In our analysis we classify each trial using the BOLD signal present seven seconds after the attention shift because of the slow characteristics of the hemodynamic response. When instead classifying an electrophysiological response obtained from implanted electrodes, a much quicker decoding will be possible. This, together with the fact that the electric signals do not suffer from the need to recover after each command, will allow BCI control at a much higher bitrate than possible with fMRI. Since the BOLD signal has been shown to map well to changes in the gamma band of ECoG [43, 89–91] the current results demonstrate the feasibility of an ECoG based BCI [43, 142, 145] controlled by visuospatial attention. While we decoded attention to four directions, the visuospatial attention control task does not have an inherent limit in number of directions. However, the more directions, the more electrodes will be needed in order to separate the responses.


6

Summary and General Discussion

Despite all the literature showing that brain-computer interfacing is achievable in a research setting, reports of paralyzed people utilizing the proposed systems are few. The performances of the BCI systems must most likely improve before they will move from research and become an assistive tool used by patients in their daily life. The majority of BCIs presented to date have been based on EEG. With the sensors placed outside the skull, this technique suffers from poor spatial resolution and low signal power. These limiting factors can be overcome by instead placing the electrodes on the cortical surface (ECoG). The opportunities for evaluating new ideas with invasive techniques are however very limited since the need for surgery excludes testing on healthy volunteers. We wished to investigate whether covert shifts of visuospatial attention could be decoded in real time and used for BCI control. With this in mind we wanted to create an environment that makes it possible to test this and other new control paradigms, to practice them with real-time performance feedback and to identify the target brain areas. The tool should be non-invasive, but still give results that are relevant for implanted technologies. Based on convincing evidence of fMRI correlating with ECoG in localization of brain functions we hypothesized that RT-fMRI could meet our needs. We implemented our RT-fMRI system on a ultra-high field (7T) MRI scanner since it could provide us with high resolution images while maintaining the high contrast-to-noise ratio that is needed to classify single trials. Using the developed system we evaluated our BCI control task, based on covert shifts of visuospatial attention, that is completely independent of any stimulation or eye movements.

In **Chapter 2** we presented results from nine healthy subjects that were scanned while instructed to move their attention to one of three spatial target regions: left, right or center. During the whole experiment, consisting of a single fMRI run, they maintained their gaze at the center. The run consisted of two parts; a localizer part in which we located the activated regions and a feedback part in which subjects were given real-time feedback based on the activity in these regions. During the feedback each image volume was classified as either left, right or center attention. The classification algorithm depended on several fixed thresholds. The values we had chosen for the thresholds turned out to be much too conservative, making most images fall into the "default" center attention class. For the last two subjects we instead used thresholds that were adapted to the individual training data. After this improvement, eight seconds into a trial of left or right attention the correct classification was 89% and 88%, respectively.

The left and right peripheral attention target areas contained checkered patterns

scrolling in opposite directions. This feature was added since we assumed it would facilitate the shifting of attention. When communicating with the subjects after the experiments they told us that they did not find the moving pattern necessary. Based on these comments and on the two successful tests we decided to use stationary attention targets in the study described in the next chapter.

Our visuospatial control paradigm has no absolute upper limit in the degrees of freedom it can provide. Building on the results from decoding the two-directions attention, in **Chapter 3** we added two additional directions, up and down. Nine healthy volunteers were scanned in a setting close to the one from Chapter 2. The main differences were that we now had four attention target areas, stationary with no scrolling checkerboards, and that we adopted a different normalization strategy. In Chapter 2 we normalized the scalar values representing the average signal within the ROIs. However, offline inspection showed that an improved performance could be achieved by first normalizing each voxel value individually and then compute the average.

When averaged over both subjects and attention directions the percentage of correctly classified trials (again eight seconds into the trial) was 79.4%. All our test trials were from one of the four peripheral directions, but we classified into five classes where the center condition was included, making the level of chance less than 25%. The strong performances across subjects added to our confidence of covert visuospatial attention being a strong target for BCI control. However, although our experiments included real-time classification with feedback, they did not yet represent a useful application.

In **Chapter 4** we decided to test the visuospatial attention task in a "real-life" BCI application, directly relevant for patients. We chose an application that takes advantage of the fact that visuospatial attention constitutes a very intuitive brain function for spatial navigation. A robot equipped with motors and a camera was navigated around a track by means of shifted attention and real-time video feedback from the camera. The video was positioned in the center of the feedback screen and surrounded by three peripheral attention target areas. In contrast to Chapters 2 and 3, where the classification was based on a univariate ROI based analysis, we now instead trained a linear support vector machine on the localizing data. The selection of which voxels to use in the training and classification was still done with the same univariate analysis as in the previous chapters. In this way we still managed to keep the experiment in a single fMRI run. Four subjects were tested in two sessions each, and they were all able to

control the robot and to reach at least three of the four targets placed in the room. The successful control shows that the foveal input from the feedback video can be distinguished from the attention to the peripheral target areas. This is very important since the central area of the visual field can then be used for feedback also in other types of BCI applications.

In **Chapter 5** we tried to bring our results closer to the ECoG modality. The four-directions attention data from Chapter 3 was reanalyzed while enforcing several spatial restrictions on the voxels available for the classifier. All the spatial restrictions applied were based on how accessible the brain areas are for implanted surface electrodes. We could show that the voxels located in brain areas suitable for surface implants provided the information necessary for successful decoding of the attention. Even when restricting the analysis to the surface of a single hemisphere, attention to the ipsilateral direction reached 70% correct classification. The classification was done using a linear support vector machine. A comparison of univariate and multivariate voxel selections showed that the latter can produce a selection resulting in a better classification.

In conclusion, our work shows that visuospatial attention can provide a novel means of controlling a BCI with multiple degrees of freedom. This control does not depend on the ability to move the eyes, and in contrast to BCI based on P300 and SSVEP it is independent of any form of visual stimulation. Our RT-fMRI system would benefit from access to a 7T compatible eye-tracker that could provide direct evidence of the absence of eye movements, without having to compare the activation pattern with a mapping of the visual cortex.

After long term paralysis or in progressed ALS, the motor cortex might no longer be capable of producing the responses utilized in a motor imagery based BCI. With the control instead based on the visual system, a BCI will not depend on an intact motor cortex.

It is unlikely that covertly directed attention to multiple directions could be decoded with EEG, especially while allowing visual feedback to be communicated at the central visual field.

Although our studies have shown that activation patterns from four directions can be distinguished, there is no reason to assume that this is the limit. It would therefore be of interest to add more attention target areas, for instance at the four diagonals making a total of eight directions.

An important question that should be addressed is, how many electrodes are necessary for a successful decoding of attention to a particular number of directions?

A component in the RT-fMRI system with large room for improvement is the voxel selection, in particular when using a multivariate classifier. A voxel selection method taking the type of classifier applied into account will be able to enhance the BCI performance. The selection algorithm should seek a group of voxels that can be covered by as few electrodes as possible and therefore the algorithm should ideally incorporate some form of cluster analysis.

Finally, the success on healthy volunteers can not automatically be extended to the conclusion that the visuospatial attention based BCI will work on LIS patients. Tests on this patient group are needed in order to show that they can orient their covert visuospatial attention as well as our test subjects.

NEDERLANDSE SAMENVATTING

Ondanks het feit dat de bestaande literatuur aangeeft dat brein-computer interactie (BCI) haalbaar is in een onderzoeksomgeving, zijn er weinig verslagen van verlamde patiënten die aangeven dergelijke systemen daadwerkelijk te gebruiken. De werking van BCI-systemen moet hoogstwaarschijnlijk nog sterk verbeterd worden voordat de BCI's de stap kunnen maken van een onderzoeksomgeving naar een daadwerkelijke ondersteuning van het dagelijkse leven van patiënten. Het grootse gedeelte van de BCI's die tot nu toe ontwikkeld zijn is gebaseerd op EEG. Aangezien hierbij de sensoren buiten de schedel geplaatst worden, heeft deze techniek last van een lage spatiële resolutie en een slechte signaal-ruisverhouding. Deze tekortkomingen kunnen teniet gedaan worden door de elektrodes direct op de cortex te plaatsen (ECOG). De mogelijkheden om nieuwe ideeën te toetsen met dergelijke invasieve technieken zijn echter beperkt vanwege de noodzaak van chirurgie, waardoor niet op gezonde vrijwilligers getest kan worden. Wij wilden onderzoeken of de hersenactiviteit bij het verschuiven van coverte visuospatiële aandacht gedecodeerd kan worden in real-time, en gebruikt kan worden voor een BCI. Met dit in het achterhoofd wilden we een omgeving creëren die het mogelijk maakt om deze en andere paradigma's te testen, ze te oefenen met real-time feedback van de prestaties, en om de belangrijkste betrokken hersengebieden te identificeren. De omgeving moet niet-invasief zijn, maar resultaten geven die van belang zijn voor implantatietechnieken. Op basis van overtuigend bewijs van overlappende fMRI en ECOG activiteit verwachtten we dat real-time fMRI (RT-fMRI) kon voldoen aan onze eisen. We hebben ons RT-fMRI systeem geïmplementeerd op een ultrahoog veld (7T) MRI scanner, omdat dit een hoge resolutie in combinatie met een hoge contrast-ruisverhouding verschafte, die nodig was om enkele trials te classificeren. Met behulp van dit systeem hebben we onze BCI-taak getest, die gebaseerd is op coverte verschuivingen van visuospatiële aandacht, en volledig onafhankelijk is van enige sensorische stimulatie of oogbewegingen.

In Hoofdstuk 2 worden de resultaten gepresenteerd van negen gezonde proefpersonen die zijn gescand terwijl ze de opdracht hadden gekregen om hun aandacht te verplaatsen naar één van drie ruimtelijke doelen op het scherm; links,rechts of mid-

den. Tijdens het hele experiment, dat bestond uit een enkele fMRI sessie, behielden zij hun blik op het centrum van het scherm. De run bestond uit twee fases: een lokalisatiefase waarin we de geactiveerde hersengebieden lokaliseerden, en een feedbackfase waarin proefpersonen real-time feedback kregen op basis van de activiteit in deze gebieden. Tijdens de feedbackfase werd ieder scanvolume geclassificeerd als aandacht links, rechts of midden. Het classificatie-algoritme was afhankelijk van enkele verschillende vaste drempelwaarden. De waarden die we hadden gekozen als drempels bleken veel te conservatief, waardoor de meeste scanvolumes werden geclassificeerd als 'centrum aandacht'. Voor de laatste twee proefpersonen gebruikten we daarom drempels die aangepast waren aan de individuele trainingsdata. Na deze aanpassing bleek dat de classificatie voor links- en rechtszijdige aandacht respectievelijk 89% en 88% correct was, op basis van het scanvolume acht seconden in het proces. De linkszijdige en rechtszijdige perifere aandachtsdoelgebieden bevatten een schaakbordpatroon dat zich verplaatste in tegengestelde richtingen. Deze eigenschap was aangebracht omdat we ervan uitgingen dat dit het verschuiven van de aandacht zou vergemakkelijken. Toen met de proefpersonen gesproken werd na het experiment vertelden ze ons dat ze het bewegende patroon niet nodig hadden om hun aandacht te richten. Gebaseerd op deze opmerkingen en twee succesvolle tests hebben we besloten om stationaire doelen te gebruiken bij de taak die in het volgende hoofdstuk beschreven wordt. Onze visuospatiële taak heeft geen absolute bovengrens in het aantal vrijheidsgraden dat het kan verschaffen. Voortbouwend op de resultaten van het decoderen van de twee aandachtsrichtingen in Hoofdstuk 3 hebben we nog twee extra richtingen toegevoegd, boven en onder. Negen gezonde vrijwilligers werden gescand met een protocol dat vrijwel gelijk was aan het protocol dat beschreven is in hoofdstuk 2. De belangrijkste verschillen waren dat er nu vier doelgebieden voor de aandacht waren, dat de doelgebieden geen bewegend schaakbordpatroon bevatten, en dat we een andere normalisatiestrategie gebruikten. In hoofdstuk 2 hadden we de scalaire waarden genormaliseerd die het gemiddelde signaal binnen de ROI representeerden. Echter, bij offline inspectie bleek dat een betere prestatie behaald kon worden door eerst de waarde van ieder individueel voxel te normaliseren en vervolgens het gemiddelde te berekenen. Bij middeling over beide proefpersonen en aandachtsrichting, was het percentage juist geclassificeerde trials (wederom geclassificeerd op het scanvolume acht seconden in het aandachtsproces) 79,4%. Al onze test trials representeerden één van vier perifere aandachtsrichtingen, maar ze werden ingedeeld in vijf klassen omdat aandacht in het centrum ook als klasse was opgenomen, waardoor het kansniveau minder dan 25% was. De hoge gemiddelde prestatie van de proefper-

sonen sterkte ons vertrouwen dat coverte visuospatiële aandacht een sterk medium is voor BCI's. Hoewel onze experimenten real-time classificatie met feedback bevatten, konden ze nog niet aangemerkt worden als een bruikbare toepassing. In Hoofdstuk 4 hebben we besloten om de visuospatiële aandachtstaak te testen in een 'real-life' BCI, waarmee het ook een voor patiënten relevante toepassing werd. We kozen voor een toepassing die gebruik maakte van het feit dat visuospatiële aandacht sterk gerelateerd is aan hersenmechanismen die betrokken zijn bij ruimtelijke navigatie. Een robot die uitgerust was met een motor en een camera werd over een circuit genavigeerd door middel van coverte visuospatiële aandacht en real-time video feedback van de camera. De video werd afgebeeld in het midden van het feedbackscherm en omgeven door drie perifere aandachtsdoelgebieden. In tegenstelling tot in de hoofdstukken 2 en 3, waarin de classificatie van de trials gebaseerd was op een univariate ROI analyse, werd nu getraind met behulp van een lineaire support vector machine. De keuze van de voxels die bij de training en classificatie gebruikt werden was nog steeds dezelfde univariate analyse als in de vorige hoofdstukken. Op deze manier kon het experiment in een enkele fMRI sessie afgenomen worden. Vier personen werden getest in twee sessies, en ze waren allemaal in staat de robot tenminste drie van de vier in de kamer geplaatste doelen te laten bereiken. De succesvolle controle laat zien dat de foveale input van de feedbackvideo kan worden onderscheiden van de aandacht voor de perifere doelgebieden. Dit is zeer belangrijk omdat het centrale deel van het gezichtsveld dan dus ook bij andere BCI toepassingen gebruikt kan worden. In Hoofdstuk 5 hebben we geprobeerd om onze resultaten dichter bij de ECOG modaliteit te brengen. De data van het experiment met de vier aandachtsrichtingen uit hoofdstuk 3 is opnieuw geanalyseerd, met oplegging van een aantal ruimtelijke beperkingen op de voxels die beschikbaar waren voor de classificatie. De toegepaste ruimtelijke beperkingen hadden betrekking op de toegankelijkheid van de hersengebieden voor geïmplanteerde oppervlakte elektroden. We konden laten zien dat de voxels in de hersengebieden die geschikt waren voor implantatie de informatie bevatte die nodig was voor een succesvolle decodering van de aandacht. Zelfs wanneer de analyse alleen gedaan werd op het oppervlak van één hemisfeer, bedroeg de correcte classificatie van aandacht in de ipsilaterale richting 70%. De classificatie werd gedaan met behulp van een lineaire support vector machine. Bij een vergelijking van univariate en multivariate classificatie bleek de laatste een betere classificatie te produceren Tot slot, ons werk laat zien dat visuospatiële aandacht gebruikt kan worden voor een nieuwe methode van BCI met meerdere vrijheidsgraden. Deze BCI is niet afhankelijk van het vermogen om de ogen te bewegen en is in tegenstelling tot BCI op P300 en SSVEP onafhankelijk

van enige vorm van visuele stimulatie. Ons RT-fMRI-systeem zou profiteren van toegang tot een 7T compatibele eye-tracker die direct bewijs zou kunnen leveren voor de afwezigheid van oogbewegingen, zonder het patroon van activiteit te moeten vergelijken met de functionele organisatie van de visuele cortex.

Na een langdurige verlamming of in vergevorderde ALS, zou de motorische cortex mogelijk niet meer in staat zijn om hersenactiviteit te genereren die toegepast kan worden bij BCI's die gebaseerd zijn op de inbeelding van motoractiviteit. Met BCI controle die gebaseerd is op het visuele systeem zal de prestatie van de BCI niet afhankelijk zijn van een intacte motorische cortex.

Het is onwaarschijnlijk dat coverte gerichte aandacht in meerdere richtingen kan worden gedecodeerd met EEG, vooral in de aanwezigheid van visuele feedback in het centrale gezichtsveld.

Hoewel onze studies hebben aangetoond dat activiteitspatronen van vier aandachtsrichtingen te onderscheiden zijn, is er geen reden om aan te nemen dat dit het hoogst haalbare is. Het zou daarom interessant zijn om meer aandachtsdoelgebieden toe te voegen, bijvoorbeeld op de vier diagonalen waardoor het totaal op acht richtingen komt.

Een belangrijke vraag die beantwoord moet worden is hoeveel elektroden nodig zijn voor een succesvolle decodering van de aandacht voor een specifiek aantal verschillende richtingen. Een aspect van de RT-fMRI waar nog veel ruimte voor verbetering is, is de voxel selectie, in het bijzonder bij het gebruik van een multivariate classificatie. Een methode van voxel selectie die rekening houdt met het type classificatie zou de prestaties van de BCI kunnen verbeteren. Het selectie algoritme moet een populatie voxels kunnen vinden die bedekt kan worden door zo weinig mogelijk elektroden en het algoritme zou dus idealiter één of andere vorm van clusteranalyse moeten bevatten.

Tenslotte kan het succes van gezonde vrijwilligers niet automatisch doorgetrokken worden naar de de conclusie dat BCI op basis van visuospatiële aandacht zal werken bij LIS patiënten. Experimenten bij deze groep patiënten zullen nodig zijn om aan te tonen dat ze hun coverte visuospatiële aandacht net zo goed kunnen oriënteren als onze gezonde proefpersonen.

Bibliography

[1] J. Tavalaro and R. Tayson. *Look up for yes*. Kodansha International, 1997.

[2] E. Smith and M. Delargy. Locked-in syndrome. *Br Med J*, 330(7488):406–409, 2005.

[3] G. Bauer, F. Gerstenbrand, and E. Rumpl. Varieties of the locked-in syndrome. *Journal of Neurology*, 221(2):77–91, 1979.

[4] C. Schnakers, A. Vanhaudenhuyse, J. Giacino, M. Ventura, M. Boly, S. Majerus, G. Moonen, and S. Laureys. Diagnostic accuracy of the vegetative and minimally conscious state: Clinical consensus versus standardized neurobehavioral assessment. *BMC Neurology*, 9(1):35, 2009.

[5] J. Mellinger, G. Schalk, C. Braun, H. Preissl, W. Rosenstiel, N. Birbaumer, and A. Kübler. An MEG-based brain-computer interface (BCI). *Neuroimage*, 36 (3):581–593, 2007.

[6] G. Sudre, L. Parkkonen, E. Bock, S. Baillet, W. Wang, and D. J. Weber. rtMEG: a real-time software interface for magnetoencephalography. *Comput Intell Neurosci*, 2011:327953, 2011.

[7] R. Sitaram, A. Caria, R. Veit, T. Gaber, G. Rota, A. Kuebler, and N. Birbaumer. fMRI brain-computer interface: A tool for neuroscientific research and treatment. *Comp Intell Neurosci*, page (10 pages), 2007.

[8] T. Moench, M. Hollmann, R. Grzeschik, C. Mueller, R. Luetzkendorf, S. Baecke, M. Luchtmann, D. Wagegg, and J. Bernarding. Real-time classification of activated brain areas for fMRI-based human-brain-interfaces. volume 6916, pages 69161R–10. SPIE, 2008.

[9] R. Sitaram, N. Weiskopf, A. Caria, R. Veit, M. Erb, and N. Birbaumer. fMRI brain-computer interfaces. *IEEE Signal Proc Mag*, 25(1):95–106, 2008.

[10] J.-H. Lee, M. Marzelli, F. A. Jolesz, and S.-S. Yoo. Automated classification of fMRI data employing trial-based imagery tasks. *Medical Image Analysis*, 13 (3):392–404, 2009.

[11] J.-H. Lee, J. Ryu, F. A. Jolesz, Z.-H. Cho, and S.-S. Yoo. Brain-machine interface via real-time fMRI: Preliminary study on thought-controlled robotic arm. *Neurosci Lett*, 450(1):1–6, 2009.

[12] N. Birbaumer, N. Ghanayim, T. Hinterberger, I. Iversen, B. Kotchoubey, A. Kubler, J. Perelmouter, E. Taub, and H. Flor. A spelling device for the paralysed. *Nature*, 398(6725):297–298, 1999.

[13] A. Kübler, N. Neumann, J. Kaiser, B. Kotchoubey, T. Hinterberger, and N. P. Birbaumer. Brain-computer communication: Self-regulation of slow cortical potentials for verbal communication. *Arch Phys Med Rehabil*, 82(11):1533–1539, 2001.

[14] T. Hinterberger, R. Veit, B. Wilhelm, N. Weiskopf, J. J. Vatine, and N. Birbaumer. Neuronal mechanisms underlying control of a brain-computer interface. *Eur J Neurosci*, 21(11):3169–81, 2005.

[15] J. R. Wolpaw, D. J. McFarland, G. W. Neat, and C. A. Forneris. An EEG-based brain-computer interface for cursor control. *Electroen Clin Neuro*, 78(3): 252–259, 1991.

[16] E. Sutter. The brain response interface: communication through visually-induced electrical brain responses. *J Microcomput Appl*, 15(1):31–45, 1992.

[17] M. Cheng, X. Gao, S. Gao, and D. Xu. Design and implementation of a brain-computer interface with high transfer rates. *IEEE Trans on Biomed Eng*, 49 (10):1181–1186, 2002.

[18] S. P. Kelly, E. C. Lalor, C. Finucane, G. McDarby, and R. B. Reilly. Visual spatial attention control in an independent brain-computer interface. *IEEE Trans on Biomed Eng*, 52(9):1588–1596, 2005.

[19] B. Z. Allison, D. J. McFarland, G. Schalk, S. D. Zheng, M. M. Jackson, and J. R. Wolpaw. Towards an independent brain-computer interface using steady state visual evoked potentials. *Clin Neurophysiol*, 119(2):399–408, 2008.

[20] M. A. Lopez-Gordo, A. Prieto, F. Pelayo, and C. Morillas. Customized stimulation enhances performance of independent binary SSVEP-BCIs. *Clin Neurophysiol*, 122(1):128–133, 2011.

[21] N. Birbaumer. Breaking the silence: Brain-computer interfaces (BCI) for communication and motor control. *Psychophysiology*, 43(6):517–532, 2006.

[22] L. A. Farwell and E. Donchin. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalogr Clin Neurophysiol*, 70(6):510–523, 1988.

[23] E. W. Sellers and E. Donchin. A P300-based brain-computer interface: initial tests by ALS patients. *Clin Neurophysiol*, 117(3):538–48, 2006.

[24] U. Hoffmann, J.-M. Vesin, T. Ebrahimi, and K. Diserens. An efficient P300-based brain-computer interface for disabled subjects. *J Neurosci Meth*, 167(1): 115–125, 2008.

[25] J. N. Mak, Y. Arbel, J. W. Minett, L. M. McCane, B. Yuksel, D. Ryan, D. Thompson, L. Bianchi, and D. Erdogmus. Optimizing the P300-based brain-computer interface: current status, limitations and future directions. *J Neural Eng*, 8(2):025003, 2011.

[26] M. Marchetti, F. Piccione, S. Silvoni, and K. Priftis. Exogenous and endogenous orienting of visuospatial attention in P300-guided brain computer interfaces: A pilot study on healthy participants. *Clin Neurophys*, 123(4):774–779, 2012.

[27] B. Z. Allison, C. Brunner, V. Kaiser, G. R. Müller-Putz, C. Neuper, and G. Pfurtscheller. Toward a hybrid brain-computer interface based on imagined movement and visual attention. *J Neural Eng*, 7(2):026007, 2010.

[28] B. Z. Allison and C. Neuper. Could anyone use a BCI?In *Brain-Computer Interfaces*. , D. S. Tan and A. Nijholt, editors, Hum-Comp Inter Series, pages 35–54. Springer London, 2010.

[29] B. Allison, T. Luth, D. Valbuena, A. Teymourian, I. Volosyak, and A. Graser. BCI demographics: How many (and what kinds of) people can use an SSVEP BCI? *IEEE Trans Neural Syst Rehabil Eng*, 18(2):107–116, 2010.

[30] A. Kübler and N. Birbaumer. Brain-computer interfaces and communication in paralysis: Extinction of goal directed thinking in completely paralysed patients? *Clin Neurophysiol*, 119(11):2658–2666, 2008.

[31] P. Brunner, S. Joshi, S. Briskin, J. R. Wolpaw, H. Bischof, and S. G. Does the 'P300' speller depend on eye gaze? *J Neur Eng*, 7(5):056013, 2010.

[32] M. S. Treder and B. Blankertz. (C)overt attention and visual speller design in an ERP-based brain-computer interface. *Behav Brain Funct*, 6(1):28, 2010.

[33] S. Kelly, E. Lalor, C. Finucane, and R. Reilly. A comparison of covert and overt attention as a control option in a steady-state visual evoked potential-based brain computer interface. In *Conf Proc IEEE Eng Med Biol Soc*, pages 4725–8, 2004.

[34] J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. L. Nicolelis. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol*, 1(2):e42, 2003.

[35] M. A. Lebedev, J. M. Carmena, J. E. O'Doherty, M. Zacksenhouse, C. S. Henriquez, J. C. Principe, and M. A. L. Nicolelis. Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *J Neurosci*, 25(19):4681–4693, 2005.

[36] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz. Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–1101, 2008.

[37] P. R. Kennedy, M. T. Kirby, M. M. Moore, B. King, and A. Mallory. Computer control using human intracortical local field potentials. *IEEE Trans Neural Syst Rehabil Eng*, 12(3):339–44, 2004.

[38] L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442 (7099):164–171, 2006.

[39] E. E. Fetz. Operant conditioning of cortical unit activity. *Science*, 163(3870): 955–958, 1969.

[40] E. C. Leuthardt, G. Schalk, J. R. Wolpaw, J. G. Ojemann, and D. W. Moran. A brain-computer interface using electrocorticographic signals in humans. *J Neural Eng*, 1(2):63–71, 2004.

[41] J. A. Wilson, E. A. Felton, P. C. Garell, G. Schalk, and J. C. Williams. ECoG factors underlying multimodal control of a brain-computer interface. *IEEE Trans Neural Syst Rehabil Eng*, 14(2):246–50, 2006.

[42] K. J. Miller, M. denNijs, P. Shenoy, J. W. Miller, R. P. Rao, and J. G. Ojemann. Real-time functional brain mapping using electrocorticography. *Neuroimage*, 37(2):504–7, 2007.

[43] M. J. Vansteensel, D. Hermes, E. J. Aarnoutse, M. G. Bleichner, G. Schalk, P. C. van Rijen, F. S. S. Leijten, and N. F. Ramsey. Brain-computer interfacing based on cognitive control. *Ann Neurol*, 67(6):809–816, 2010.

[44] E. Margalit, J. D. Weiland, R. E. Clatterbuck, G. Y. Fujii, M. Maia, M. Tameesh, G. Torres, S. A. D'Anna, S. Desai, D. V. Piyathaisere, A. Olivi, J. de Juan, E, and M. S. Humayun. Visual and electrical evoked response recorded from subdural electrodes implanted above the visual cortex in normal dogs under two methods of anesthesia. *J Neurosci Methods*, 123(2):129–37, 2003.

[45] E. C. Leuthardt, G. Schalk, J. Roland, A. Rouse, and D. W. Moran. Evolution of brain-computer interfaces: going beyond classic motor physiology. *Neurosurg Focus*, 27(1):E4, 2009.

[46] A. Gunduz, P. Brunner, A. Daitch, E. C. Leuthardt, A. L. Ritaccio, B. Pesaran, and G. Schalk. Neural correlates of visual spatial attention in electrocorticographic (ECoG) signals in humans. *Frontiers in Human Neuroscience*, 5, 2011.

[47] B. A. Wandell, S. O. Dumoulin, and A. A. Brewer. Visual field maps in human cortex. *Neuron*, 56(2):366–383, 2007.

[48] S. Kastner, P. De Weerd, R. Desimone, and L. G. Ungerleider. Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science*, 282(5386):108, 1998.

[49] J. A. Brefczynski and E. A. DeYoe. A physiological correlate of the 'spotlight' of visual attention. *Nat Neurosci*, 2(4):370–374, 1999.

[50] N. G. Müller and D. Ebeling. Attention-modulated activity in visual cortex - more than a simple 'spotlight'. *Neuroimage*, 40(2):818–827, 2008.

[51] R. Datta and E. A. DeYoe. I know where you are secretly attending! the topography of human visual attention revealed with fMRI. *Vision Res*, 49(10): 1037–1044, 2009.

[52] J. A. Brefczynski-Lewis, R. Datta, J. W. Lewis, and E. A. DeYoe. The topography of visuospatial attention as revealed by a novel visual field mapping technique. *J Cognitive Neurosci*, 21(7):1447–1460, 2009.

[53] D. C. Somers, A. M. Dale, A. E. Seiffert, and R. B. H. Tootell. Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proc Natl Acad Sci USA*, 96(4):1663–1668, 1999.

[54] L. Pessoa, S. Kastner, and L. G. Ungerleider. Neuroimaging studies of attention: From modulation of sensory processing to top-down control. *J Neurosci*, 23(10):3990–3998, 2003.

[55] M. A. Silver, D. Ress, and D. J. Heeger. Neural correlates of sustained spatial attention in human early visual cortex. *J Neurophysiol*, 97(1):229–237, 2007.

[56] L. Carroll. *Sylvie and Bruno*. Macmillan and Co, 1889.

[57] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci USA*, 87(24):9868–9872, 1990.

[58] P. A. Bandettini, E. C. Wong, R. S. Hinks, R. S. Tikofsky, and J. S. Hyde. Time course EPI of human brain function during task activation. *Magnet Reson Med*, 25(2):390–397, 1992.

[59] K. K. Kwong, J. W. Belliveau, D. A. Chesler, I. E. Goldberg, R. M. Weisskoff, B. P. Poncelet, D. N. Kennedy, B. E. Hoppel, M. S. Cohen, and R. Turner. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc Natl Acad Sci USA*, 89(12):5675–5679, 1992.

[60] S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci USA*, 89(13):5951–5955, 1992.

[61] M. E. Raichle. A brief history of human brain mapping. *Trends Neurosci*, 32 (2):118–26, 2009.

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

[63] L. Pauling and C. D. Coryell. The magnetic properties and structure of hemoglobin, oxyhemoglobin and carbonmonoxyhemoglobin. *Proc Natl Acad Sci USA*, 22(4):210–6, 1936.

[64] M. S. Cohen and S. Y. Bookheimer. Localization of brain function using magnetic resonance imaging. *Trends Neurosci*, 17(7):268–277, 1994.

[65] S. Ogawa, R. S. Menon, S.-G. Kim, and K. Ugurbil. On the characteristics of functional magnetic resonance imaging of the brain. *Annu Rev Biophys Biomol Struct*, 27(1):447–474, 1998.

[66] M. A. Lindquist, J. Meng Loh, L. Y. Atlas, and T. D. Wager. Modeling the hemodynamic response function in fMRI: efficiency, bias and mis-modeling. *Neuroimage*, 45(1 Suppl):S187–98, 2009.

[67] W. van der Zwaag, S. Francis, K. Head, A. Peters, P. Gowland, P. Morris, and R. Bowtell. fMRI at 1.5, 3 and 7 T: Characterising BOLD signal changes. *Neuroimage*, 47(4):1425–1434, 2009.

[68] R. Beisteiner, S. Robinson, M. Wurnig, M. Hilbert, K. Merksa, J. Rath, I. Höllinger, N. Klinger, C. Marosi, S. Trattnig, and A. Geißler. Clinical fMRI: Evidence for a 7T benefit over 3T. *Neuroimage*, 57(3):1015–1021, 2011.

[69] J. C. W. Siero, N. Petridou, H. Hoogduin, P. R. Luijten, and N. F. Ramsey. Cortical depth-dependent temporal dynamics of the BOLD response in the human brain. *J Cereb Blood Flow Metab*, 31(10):1999–2008, 2011.

[70] K. J. Friston, A. P. Holmes, K. J. Worsley, J. P. Poline, C. D. Frith, and R. S. J. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Hum Brain Map*, 2(4):189–210, 1994.

[71] R. W. Cox, A. Jesmanowicz, and J. S. Hyde. Real-time functional magnetic resonance imaging. *Magn Reson Med*, 33(2):230–6, 1995.

[72] A. Caria, R. Sitaram, and N. Birbaumer. Real-time fMRI: A tool for local brain regulation. *The Neuroscientist*, In Press.

[73] G. Fernandez, A. de Greiff, J. von Oertzen, M. Reuber, S. Lun, P. Klaver, J. Ruhlmann, J. Reul, and C. E. Elger. Language mapping in less than 15 minutes: Real-time functional MRI during routine clinical investigation. *Neuroimage*, 14(3):585–594, 2001.

[74] R. C. deCharms, F. Maeda, G. H. Glover, D. Ludlow, J. M. Pauly, D. Soneji, J. D. Gabrieli, and S. C. Mackey. Control over brain activation and pain learned by using real-time functional MRI. *Proc Natl Acad Sci USA*, 102(51):18626–31, 2005.

[75] S. Haller, N. Birbaumer, and R. Veit. Real-time fMRI feedback training may improve chronic tinnitus. *Eur Radiol*, 20(3):696–703, 2010.

[76] P. Van Gelderen, N. F. Ramsey, G. Liu, J. H. Duyn, J. A. Frank, D. R. Weinberger, and C. T. W. Moonen. Three-dimensional functional magnetic resonance imaging of human brain on a clinical 1.5-T scanner. *Proc Natl Acad Sci USA*, 92(15):6906–6910, 1995.

[77] A. M. Smith, B. K. Lewis, U. E. Ruttimann, F. Q. Ye, T. M. Sinnwell, Y. Yang, J. H. Duyn, and J. A. Frank. Investigation of low frequency drift in fMRI signal. *Neuroimage*, 9(5):526–533, 1999.

[78] J. Tanabe, D. Miller, J. Tregellas, R. Freedman, and F. G. Meyer. Comparison of detrending methods for optimal fMRI preprocessing. *Neuroimage*, 15(4): 902–907, 2002.

[79] J.-D. Haynes and G. Rees. Decoding mental states from brain activity in humans. *Nat Rev Neurosci*, 7(7):523–534, 2006.

[80] M. Misaki, Y. Kim, P. A. Bandettini, and N. Kriegeskorte. Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *Neuroimage*, 53(1):103–118, 2010.

[81] S. M. LaConte, S. J. Peltier, and X. P. Hu. Real-time fMRI using brain-state classification. *Hum Brain Mapp*, 28(10):1033–1044, 2007.

[82] M. Hollmann, J. W. Rieger, S. Baecke, R. Lützkendorf, C. Müller, D. Adolf, and J. Bernarding. Predicting decisions in human social interactions using real-time fMRI and pattern classification. *PLoS ONE*, 6(10):e25304, 2011.

[83] R. Sitaram, S. Lee, S. Ruiz, M. Rana, R. Veit, and N. Birbaumer. Real-time support vector classification and feedback of multiple emotional brain states. *Neuroimage*, 56(2):753–65, 2011.

[84] S. M. LaConte. Decoding fMRI brain states in real-time. *Neuroimage*, 56(2): 440–454, 2011.

[85] F. De Martino, G. Valente, N. Staeren, J. Ashburner, R. Goebel, and E. Formisano. Combining multivariate voxel selection and support vector machines for mapping and classification of fMRI spatial patterns. *Neuroimage*, 43 (1):44–58, 2008.

[86] C. Smyser, T. J. Grabowski, R. J. Frank, J. W. Haller, and L. Bolinger. Real-time multiple linear regression for fMRI supported by time-aware acquisition and processing. *Magn Reson Med*, 45(2):289–98, 2001.

[87] E. Bagarinao, K. Matsuo, T. Nakai, and S. Sato. Estimation of general linear model coefficients for real-time application. *Neuroimage*, 19(2):422–429, 2003.

[88] O. Hinds, S. Ghosh, T. W. Thompson, J. J. Yoo, S. Whitfield-Gabrieli, C. Triantafyllou, and J. D. E. Gabrieli. Computing moment-to-moment BOLD activation for real-time neurofeedback. *Neuroimage*, 54(1):361–368, 2011.

[89] J. P. Lachaux, P. Fonlupt, P. Kahane, L. Minotti, D. Hoffmann, O. Bertrand, and M. Baciu. Relationship between task-related gamma oscillations and BOLD signal: New insights from combined fMRI and intracranial EEG. *Hum Brain Mapp*, 28(12):1368–1375, 2007.

[90] B. J. He, A. Z. Snyder, J. M. Zempel, M. D. Smyth, and M. E. Raichle. Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. *Proc Natl Acad Sci USA*, 105(41):16039–16044, 2008.

[91] D. Hermes, K. J. Miller, M. J. Vansteensel, E. J. Aarnoutse, F. S. S. Leijten, and N. F. Ramsey. Neurophysiologic correlates of fMRI in human motor cortex. *Hum Brain Map*, In Press.

[92] J. R. Wolpaw, N. Birbaumer, D. J. McFarland, G. Pfurtscheller, and T. M. Vaughan. Brain-computer interfaces for communication and control. *Clin Neurophysiol*, 113(6):767–791, 2002.

[93] N. Birbaumer and L. G. Cohen. Brain-computer interfaces: communication and restoration of movement in paralysis. *J Physiol*, 579(3):621–36, 2007.

[94] M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. P. Donoghue. Brain-machine interface: Instant neural control of a movement signal. *Nature*, 416(6877):141–142, 2002.

[95] D. M. Taylor, S. I. H. Tillery, and A. B. Schwartz. Direct cortical control of 3D neuroprosthetic devices. *Science*, 296(5574):1829–1832, 2002.

[96] M. A. L. Nicolelis. Brain-machine interfaces to restore motor function and probe neural circuits. *Nat Rev Neurosci*, 4(5):417–422, 2003.

[97] B. Graimann, J. E. Huggins, S. P. Levine, and G. Pfurtscheller. Toward a direct brain interface based on human subdural recordings and wavelet-packet analysis. *IEEE Trans on Biomed Eng*, 51(6):954–962, 2004.

[98] J. C. Sanchez, A. Gunduz, P. R. Carney, and J. C. Principe. Extraction and localization of mesoscopic motor control signals for human ECoG neuroprosthetics. *J Neurosci Meth*, 167(1):63–81, 2008.

[99] P. Shenoy, K. J. Miller, J. G. Ojemann, and R. P. Rao. Generalized features for electrocorticographic BCIs. *IEEE Trans Biomed Eng*, 55(1):273–80, 2008.

[100] J. Kalcher, D. Flotzinger, C. Neuper, S. Gölly, and G. Pfurtscheller. Graz brain-computer interface II: towards communication between humans and computers based on online classification of three different EEG patterns. *Med Biol Eng Comput*, 34(5):382–388, 1996.

[101] C. Guger, G. Edlinger, W. Harkam, I. Niedermayer, and G. Pfurtscheller. How many people are able to operate an EEG-based brain-computer interface (BCI)? *IEEE Trans Neural Syst Rehabil Eng*, 11(2):145–147, 2003.

[102] C. Guger, S. Daban, E. Sellers, C. Holzner, G. Krausz, R. Carabalona, F. Gramatica, and G. Edlinger. How many people are able to control a P300-based brain-computer interface (BCI)? *Neurosci Lett*, 462(1):94–98, 2009.

[103] B. Blankertz, C. Sannelli, S. Halder, E. M. Hammer, A. Kübler, K.-R. Müller, G. Curio, and T. Dickhaus. Neurophysiological predictor of SMR-based BCI performance. *Neuroimage*, 51(4):1303–1309, 2010.

[104] E. I. Knudsen. Fundamental components of attention. *Annu Rev Neurosci*, 30 (1):57–78, 2007.

[105] M. I. Posner. Orienting of attention. *Q J Exp Psychol*, 32(1):3–25, 1980.

[106] M. van Gerven, A. Bahramisharif, T. Heskes, and O. Jensen. Selecting features for BCI control based on a covert spatial attention paradigm. *Neural Networks*, 22(9):1271–1277, 2009.

[107] S. P. Kelly, E. Lalor, R. B. Reilly, and J. J. Foxe. Independent brain computer interface control using visual spatial attention-dependent modulations of parieto-occipital alpha. In *2nd Int IEEE EMBS Conf on Neural Eng*, pages 667–670, 2005.

[108] N. Weiskopf, R. Veit, M. Erb, K. Mathiak, W. Grodd, R. Goebel, and N. Birbaumer. Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. *Neuroimage*, 19(3):577–86, 2003.

[109] S. S. Yoo, T. Fairneny, N. K. Chen, S. E. Choo, L. P. Panych, H. Park, S. Y. Lee, and F. A. Jolesz. Brain-computer interface using fMRI: spatial navigation by thoughts. *Neuroreport*, 15(10):1591–1595, 2004.

[110] R. C. deCharms. Applications of real-time fMRI. *Nat Rev Neurosci*, 9(9): 720–729, 2008.

[111] S. Klein, M. Staring, and J. P. W. Pluim. Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines. *IEEE Trans Image Process*, 16(12):2879–2890, 2007.

[112] S. Schwartz, P. Vuilleumier, C. Hutton, A. Maravita, R. J. Dolan, and J. Driver. Attentional load and sensory competition in human vision: Modulation of fMRI responses by load at fixation during task-irrelevant stimulation in the peripheral visual field. *Cerebral Cortex*, 15(6):770–786, 2005.

[113] N. G. Müller, M. Mollenhauer, A. Rösler, and A. Kleinschmidt. The attentional field has a mexican hat distribution. *Vision Res*, 45(9):1129–1137, 2005.

[114] M. P. Tarvainen, P. O. Ranta-aho, and P. A. Karjalainen. An advanced detrending method with application to HRV analysis. *IEEE Trans Biomed Eng*, 49(2): 172–175, 2002.

[115] N. F. Ramsey, M. P. van de Heuvel, K. H. Kho, and F. S. S. Leijten. Towards human BCI applications based on cognitive brain systems: an investigation of neural signals recorded from the dorsolateral prefrontal cortex. *IEEE Trans Neural Syst Rehabil Eng*, 14(2):214–217, 2006.

[116] J. R. Wolpaw, H. Ramoser, D. J. McFarland, and G. Pfurtscheller. EEG-based communication: improved accuracy by response verification. *IEEE Trans Rehabil Eng*, 6(3):326–333, 1998.

[117] L. Heinemann, A. Kleinschmidt, and N. G. Müller. Exploring BOLD changes during spatial attention in non-stimulated visual cortex. *PLoS ONE*, 4(5):e5560, 2009.

[118] S. O. Dumoulin, R. G. Bittar, N. J. Kabani, C. L. Baker, G. Le Goualher, G. B. Pike, and A. C. Evans. A new anatomical landmark for reliable identification of human area V5/MT: a quantitative analysis of sulcal patterning. *Cereb Cortex*, 10(5):454–463, 2000.

[119] R. F. Dougherty, V. M. Koch, A. A. Brewer, B. Fischer, J. Modersitzki, and B. A. Wandell. Visual field representations and locations of visual areas V1/2/3 in human visual cortex. *J Vis*, 3(10), 2003.

[120] F. Lotte, M. Congedo, A. Lécuyer, F. Lamarche, and B. Arnaldi. A review of classification algorithms for EEG-based brain-computer interfaces. *J Neural Eng*, 4(2):R1, 2007.

[121] F. Popescu, B. Blankertz, and K.-R. Müller. Computational challenges for non-invasive brain computer interfaces. *IEEE Intell Syst*, 23(3):78–79, 2008.

[122] A. Kübler, F. Nijboer, J. Mellinger, T. M. Vaughan, H. Pawelzik, G. Schalk, D. J. McFarland, N. Birbaumer, and J. R. Wolpaw. Patients with als can use sensorimotor rhythms to operate a brain-computer interface. *Neurology*, 64 (10):1775–1777, 2005.

[123] T. Siman-Tov, A. Mendelsohn, T. Schonberg, G. Avidan, I. Podlipsky, L. Pessoa, N. Gadoth, L. G. Ungerleider, and T. Hendler. Bihemispheric leftward bias in a visuospatial attention-related network. *J Neurosci*, 27(42):11271–11278, 2007.

[124] J. Munneke, D. J. Heslenfeld, and J. Theeuwes. Directing attention to a location in space results in retinotopic activation in primary visual cortex. *Brain Res*, 1222:184–191, 2008.

[125] M. Treder, A. Bahramisharif, N. Schmidt, M. van Gerven, and B. Blankertz. Brain-computer interfacing using modulations of alpha activity induced by covert shifts of attention. *J NeuroEng Rehabil*, 8(1):24, 2011.

[126] M. van Gerven and O. Jensen. Attention modulations of posterior alpha as a control signal for two-dimensional brain-computer interfaces. *J Neurosci Meth*, 179(1):78–84, 2009.

[127] A. Bahramisharif, M. Van Gerven, T. Heskes, and O. Jensen. Covert attention allows for continuous control of brain-computer interfaces. *Eur J Neurosci*, 31 (8):1501–1508, 2010.

[128] E. Yacoub, A. Shmuel, J. Pfeuffer, P.-F. Van De Moortele, G. Adriany, P. Andersen, J. T. Vaughan, H. Merkle, K. Ugurbil, and X. Hu. Imaging brain function in humans at 7 Tesla. *Magn Res Med*, 45(4):588–594, 2001.

[129] S. A. Engel, D. E. Rumelhart, B. A. Wandell, A. T. Lee, G. H. Glover, E.-J. Chichilnisky, and M. N. Shadlen. fMRI of human visual cortex. *Nature*, 369 (6481):525–525, 1994.

[130] M. I. Sereno, A. M. Dale, J. B. Reppas, K. K. Kwong, J. W. Belliveau, T. J. Brady, B. R. Rosen, and R. B. Tootell. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, 268(5212): 889–893, 1995.

[131] D. Van Essen, J. Dickson, J. Harwell, D. Hanlon, C. Anderson, and H. Drury. An integrated software system for surface-based analyses of cerebral cortex. *J Am Med Inform Assoc*, 8(5):443–459, 2001.

[132] P. R. Roelfsema. Attention-voluntary control of brain cells. *Science*, 332(6037): 1512–1513, 2011.

[133] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM T Int Syst Tech*, 2(3):1–27, 2011.

[134] S. LaConte, S. Strother, V. Cherkassky, J. Anderson, and X. Hu. Support vector machines for temporal classification of block design fMRI data. *Neuroimage*, 26(2):317–329, 2005.

[135] S. L. Shishkin, I. P. Ganin, I. A. Basyul, A. Y. Zhigalov, and A. Y. Kaplan. N1 wave in the P300 BCI is not sensitive to the physical characteristics of stimuli. *J Integr Neurosci*, 8(4):471–85, 2009.

[136] L. Bianchi, S. Sami, A. Hillebrand, I. P. Fawcett, L. R. Quitadamo, and S. Seri. Which physiological components are more suitable for visual ERP based brain-computer interface? A preliminary MEG/EEG study. *Brain Topogr*, 23(2): 180–5, 2010.

[137] R. J. Perry and S. Zeki. The neurology of saccades and covert shifts in spatial attention. *Brain*, 123(11):2273–2288, 2000.

[138] H. Yamamoto, M. Fukunaga, S. Takahashi, H. Mano, C. Tanaka, M. Umeda, and Y. Ejima. Inconsistency and uncertainty of the human visual area loci following surface-based registration: Probability and entropy maps. *Hum Brain Map*, 33(1):121–129, 2012.

[139] M. S. Treder, N. M. Schmidt, and B. Blankertz. Gaze-independent brain-computer interfaces based on covert attention and feature attention. *J Neural Eng*, 8(6), 2011.

[140] P. R. Kennedy and R. A. E. Bakay. Restoration of neural output from a paralyzed patient by a direct brain connection. *Neuroreport*, 9(8):1707–1711, 1998.

[141] P. R. Kennedy, R. A. E. Bakay, M. M. Moore, K. Adams, and J. Goldwaithe. Direct control of a computer from the human central nervous system. *IEEE Trans Rehabil Eng*, 8(2):198–202, 2000.

[142] T. Yanagisawa, M. Hirata, Y. Saitoh, T. Goto, H. Kishima, R. Fukuma, H. Yokoi, Y. Kamitani, and T. Yoshimine. Real-time control of a prosthetic hand using human electrocorticography signals. *J Neurosurg*, 114(6):1715–1722, 2011.

[143] M. Hanke, Y. Halchenko, P. Sederberg, S. Hanson, J. Haxby, and S. Pollmann. PyMVPA: a Python toolbox for multivariate pattern analysis of fMRI data. *Neuroinform*, 7(1):37–53, 2009.

[144] I. Guyon, J. Weston, S. Barnhill, and V. Vapnik. Gene selection for cancer classification using support vector machines. *Mach Learn*, 46(1):389–422, 2002.

[145] G. Schalk and E. Leuthardt. Brain-computer interfaces using electrocorticographic (ECoG) signals. *IEEE Rev Biomed Eng*, 4:140–54, 2011.

Acknowledgments

Finally ....
No more time spent doubting that all the hard work will result in this book. The process has been far from smooth, but that does not mean it has been without joy. I have had a great time and learned so much during my PhD. One thing is certain, I would not have managed without the help and support from a large number of people. They provided me with everything from knowledge about neurosurgical procedures and MR sequences, to a beer or a slap in the face while telling me to focus. Without the contributions from all these people my work would not have been possible.

First of all, I would like to thank Max for giving me the opportunity to work in the inspiring environment of the Image Sciences Institute. Thank you for the freedom you have given me in my project and for not giving up on me when things seemed to be going against me. You have stayed incredibly supportive during the whole process.

Nick Ramsey, as the man behind the BCI project you are the one who made it possible for me to work in this fantastic field. Your enthusiasm and knowledge has been truly inspiring. You were always available for moral and technical support and have stayed positive during my entire project. In fact, often you had to sell my work to me instead of the other way around. I really hope there will be possibilities for further collaborations in the future.

Josien Pluim, many times I have entered your room frustrated and left refueled with hope, whether or not the source of my frustration was work related. Though my work was slightly outside your field, you have always been available for valuable discussions and you have been a great guide in steering my project when it was going in the wrong direction.

To all my colleagues at ISI, both current and former, thank you for making the group the inspiring place to work in that it is. Adriënne, Casper, Chris, Gerrit, Hans, Hendrik, Hugo, Ivana, Jaco, Jeroen B, Jeroen S, Joffrey, Keelin, Koen, Maartje, Marijn, Marius, Martijn, Mitko, Nynke, Paul, Peter, Rashindra, Robin, Roel, Sandra, Sara, Sascha, Stefan, Vincent, Wilbert and the rest of you, thank you all for interesting discussions, nice coffee breaks and just generally good times. However, some names deserve a special mention.
When I had medically related questions I could always count on Jeroen and Casper.

Thank you for this, and for the great time I had sharing a room with you.
Speaking of room mates (I had a few), Sara, Martijn, Jeroen, Casper, Sandra, Nynke, Aldemar and Mathijs: you always made sure it was a pleasure entering the office.
Martijn, you convinced me to not give up, and for this I owe you big time. If it was not for you, instead of asking myself 'Did I really finish?', I would be asking someone else 'Do you want fries with that?'. Sara, I never laughed as much at work as when we shared an office (even excluding the riverdance). I hope I get to visit you and Martijn in Boston soon. Stefan, thank you for all your effort helping me get the real-time registration working. You are truly one of the kindest and most helpful persons I have met. I'm sure we'll keep meeting up for concerts and beer. When I had MRI related problems, I was in heaven. With the combined knowledge of Wilbert, Natalia, Hans, Fredy, Jaco and Jeroen it was simply impossible to not end up with a great solution. Gerard, thanks for all your help with computer related issues, and for being a good source of cables. Renée, Marjan and Jacqueline, thank you for making sure things run smoothly and for always helping out with a smile.

While I belonged to ISI, my work was done mainly together with the BCI group. Working with the BCI project has been absolutely fantastic, much due to the people involved. Jeroen, thanks for the many hours you spent at the scanner during my experiments. Martin, Zac, Mariska, Erik and Dora you were always available for helpful insights into the world of electrophysiology. I also want to thank you all for the help with the ECoG experiments and Martin and Erik for joining me to Graz for the fNIRS experiments. Aldemar, I could always count on you to share my love for a good equation. Many thanks also to Annemiek for providing some structure to it all.

Mathijs and Gert, as fellow fMRIers you were very useful, both inside and outside the bore. Mathijs, thanks for helping me inflate and flatten brains (now, thats a sentence I never thought I would write). You were also a good compass when I lost my orientation in the visual field maps. Gert, we had many interesting discussions and you made sure I was on the correct dosage of coffee and beer.

I also want to give credit to the flow managers Estrella, Lotte and David who somehow made it fun to spend a Friday evening inside a dark bunker. On that note, I should also acknowledge all the volunteers that allowed me to borrow their brains for an hour or two.

The reason that I moved to the Netherlands in the first place was to do a Masters project at the VU for Jan de Munck. Jan, it was in large part thanks to you that I got the chance to work at ISI and I want to thank you for this. I had an amazing time in your

lab and a couple of my colleagues during those days turned in to close friends. Sandra (and your incredible family!), Rui and Teresa, thank you so much for your welcoming ways and for making sure I have a place to go when I need to clear my head (and eat Pastéis de Belém).

Even the best job in the world would be unbearable without friends supporting a quality life outside work. If someone reading this should need a good friend, look for one of these names: Mel, Roy, Ane, Jennifer, Susie, Patrick, Margriet, Markus, Miriam, Linus, Nathalie, Erik, Ricko, David, Maria, Alev, Shadi. They know their stuff.

A huge thanks to my family for being a constant source of all possible kinds of support. The same goes for "my other family" that Mia kindly provided me with.

And finally, Mia, you are the best collection of molecules that ever clustered. Between you and me, all the people mentioned above are terrific, but nothing is truly great without you.

Publications

Papers in international journals

- P. Andersson, J.P.W. Pluim, J.C.W. Siero, S. Klein, M.A. Viergever and N.F. Ramsey, "Real-Time Decoding of Brain Responses to Visuospatial Attention Using 7T fMRI", *PLoS ONE*, 6(11):e27638, 2011
- P. Andersson, N.F. Ramsey, M. Raemaekers, M.A. Viergever and J.P.W. Pluim, "Real-Time Decoding of Direction of Covert Visuospatial Attention", *Journal of Neural Engineering*, (In Press)
- P. Andersson, J.P.W. Pluim, M.A. Viergever and N.F. Ramsey, "Navigation of a Telepresence Robot via Covert Visuospatial Attention", (Submitted)
- P. Andersson, N.F. Ramsey, M.A. Viergever and J.P.W. Pluim, "7T fMRI for Identification of Cortical Target Regions for Intracranial Brain-Computer Interfacing based on Covert Visuospatial Attention", (In Preparation)

Papers in conference proceedings

- S. Klein, M. Staring, P. Andersson and J.P. Pluim, "Preconditioned Stochastic Gradient Descent Optimisation for Monomodal Image Registration", *Medical Image Computing and Computer-Assisted Intervention 2011*, 14(Pt 2):549-56
- P. Andersson, N.F. Ramsey, J.P.W. Pluim and M.A. Viergever, "BCI Control Using 4 Direction Spatial Visual Attention and Real-Time fMRI at 7T", *Proc. of IEEE Engineering in Medicine and Biology Society 2010*, 4221-25, 2010
- N.F. Ramsey, J.C.W. Siero, J.P. Pluim, M.A. Viergever and P. Andersson, "High-Field fMRI for Localization of Electrode Implant Sites for BCI: A Study on Visual Attention", *BCI Meeting 2010*, 2010
- P. Andersson, M.A. Viergever, J.P.W. Pluim, N.F. Ramsey, J.C.W Siero, "fMRI Based BCI Control using Spatial Visual Attention at 7T", *IEEE EMBS Conference on Neural Engineering 2009*, 444-46, 2009

Abstracts

- P. Andersson, J.C.W. Siero, J.P.W. Pluim, M.A. Viergever and N.F. Ramsey, "Visual Attention for Brain-Computer Interface: Towards Using 7T fMRI to Localize Electrode Implant Sites", *International Society for Magnetic Resonance in Medicine*, 2010
- P. Andersson, J.P.W. Pluim, M.A. Viergever and N.F. Ramsey, "Attentional Modulation of Visual Cortex BOLD Activity as a Basis for Brain-Computer Interface Systems", *Society for Neuroscience 2010*, 2010
- P. Andersson, J.C.W. Siero, J.P.W. Pluim, M.A. Viergever and N.F. Ramsey, "Real-Time Single-Trial BOLD Response Detection for Visual Attention at 7T", *International Society for Magnetic Resonance in Medicine*, 2009