Using a Support-Vector Machine for the analysis of fMRI images in a tactile frequency-discrimination task.

J.B.C. Marsman

1204440

Master's Thesis in Artificial Intelligence
April 24, 2006

External Supervisor: Dr. L.M. Parkes
Internal Supervisor: Prof. Dr. L.R.B. Schomaker

University of Groningen
Department of Behavioral and Social Sciences
Grote Kruisstraat 1-2
The Netherlands

957

Abstract

Traditionally, fMRI analysis uses a univariate method for statistical testing, such as the General Linear Model. In the past few years multivariate approaches have been developed also taking spatial relations between voxels into account, including Support-Vector Machines (SVMs). An SVM is a binary classifier which calculates the optimal hyperplane (decision boundary) between two classes in a training set. This hyperplane is then used to predict the class of the items in a test set.

The goal of this study was to find correlates between human discrimination ability and the performance of the SVM on the captured fMRI images from this discrimination task. The task was to discriminate between different vibrotactile frequencies, where it is still an open question as to whether neural firing rate or synchronicity is used. fMRI signal intensity changes as a function of firing rate, but not synchronicity, and so could inform this question.

We conducted an fMRI experiment where vibrotactile stimuli were applied to the right index finger, consisting of a reference frequency and a stimulus frequency. After each pair of stimuli, the subject had to respond with either higher or lower, using a button box held in the left hand. Beside standard fMRI preprocessing and stimulus-related averaging over each session, we used either Singular Value Decomposition or feature selection methods to decrease computational time.

The results led to the conclusion that frequency is not spatially represented. This still leaves rate encoding and periodicity encoding to be possible representations of frequency.

Contents

Chapter 1

Introduction

1.1 Introduction

Within the field of neuroscience one of the aims is to map the cognitive state with the underlying neural activity of the human brain. This is commonly referred to as finding 'neural correlates' of cognitive states. With this graduation project I have been trying to find a neural correlate in the somatosensory cortex regarding the representation of vibration, because it is still unknown how vibrotactile stimuli are represented in the somatosensory cortex. The somatosensory cortex is the part of the brain which carries out all lower level cognitive tasks regarding somatic sensation.

To investigate how these stimuli are represented, I have conducted a vibrotactile discrimination task inside a Magnetic Resonance scanner. Showing how vibrotactile stimuli are processed provides new insight into the understanding of the working of the brain. Using methods from pattern recognition, I have tried to find a correlation between the frequency of the represented stimulus and the activity in the somatosensory cortex during this stimulus. The activity of the brain is captured using the MR scanner.

This research project integrates methods from artificial intelligence with the neuroscience of sensory perception. It has been performed at the Magnetic Resonance Imaging and Research Analysis Centre (MARIARC) of the University of Liverpool in the United Kingdom.

1.2 Research objective

The aim of this project is to find correlates between on the one hand the subjects performance on a perceptual discrimination task and on the other hand the classification performance of statistical discrimination (support-vector machines) of the images taken with the MR scanner during this task.

To investigate this possible relationship, the following main question has been stated and split up into three subquestions.

Is the support-vector machine able to provide information about the performance of the subject?

with the following sub questions:

1. What is the general classifier performance on large sets of brain images collected either during the application of vibration stimulation or during rest?
2. What is the human performance on the vibrotactile discrimination in comparison to classifier performance on the discrimination of cognitive states?
3. What do learning trends of the vibrotactile discrimination task by humans look like?

Making hypotheses on this matter, one may expect that as the difficulty of the discrimination task increases, the human performance will decrease. This decrease has to do with the tactile acuity of human perception of vibration.

With this decrease the performance of classifying the corresponding brain states by Support-Vector Machines is also likely to decrease. Reflecting that this neural decision is represented in the activity of the brain. This would infer a positive correlation, which will prove the possibility to classify vibrational somatosensory stimuli. Since the subjects performance is likely to contain a small learning trend over time, a consequent hypothesis is that the performance of the Support-Vector Machine for the datasets corresponding to the first experiment will be lower than those corresponding to the final datasets.

1.3 Content Overview

This thesis can be divided into three main topics:

- Neuroscientific background information (Chapters 2 - 4).
  The next chapter contains a detailed description of the mechanisms involved in tactile sensory information processing, followed by basic principles of Magnetic Resonance Imaging and methods of analysis in neuroscience.
- Description of Support-Vector Machines (Chapter 5).
  This chapter will discuss the statistical idea of support-vector machines and its limitations. Furthermore an overview will be given of support-vector machines applied in functional MRI studies.
- Description of the conducted experiment (Chapters 6 - 9).
  The actual research and performed analyses will be discussed in chapter six, followed by results, discussion and conclusion.

Chapter 2

Somatosensory Information Processing

In order to investigate the representation of vibrotactile stimuli, a clear overview has to be made about the way the brain processes vibrotactile stimuli. The following chapter provides an overview of the somatosensory system which is used by humans to have tactile sensations. Our somatosensory system provides the capability of having a continuous representation of the external world. (Zigmond et al., 1999).

From the location where the stimulus is presented, the vibration is translated into a neural code by the receptors and transported by the afferents (see Figure 2.1(b)). This code eventually arrives in the brain which performs the requested comparison. In respect of the experiment a certain motor action will follow in correspondence with the outcome of the comparison. From the research of Penfield and Rasmussen in 1950, we know that the sensory information is somatotopically organised along the central sulcus (see Figure 2.1(a)). This knowledge enables us to zoom in on a specific region where the somatosensory information is being processed, the somatosensory cortex. The somatosensory cortex is located in the postcentral gyrus.

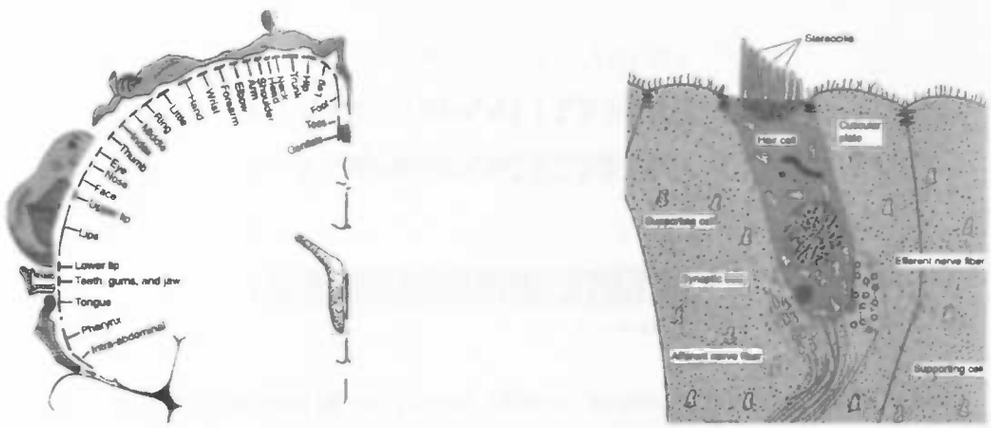

(a) Penfield's topological map of the central gyrus (Zigmond et al., 1999)

(b) Receptors and afferents used in transduction of tactile stimulation (Zigmond et al., 1999)

Figure 2.1: Components of the Somatosensory system

The somatosensory system can be functionally divided into three categories:

Exteroceptive functions This class of functions concerns the sense of touch (mechanoreception), pain (nociception) and temperature (thermoreception). Since we are only using mechanoreception in this project this chapter will only explore the properties of vibrational mechanoreception.

Proprioceptive functions These functions are used when feeling position and movement. They are driven by muscular input and provide information on relative directions, speed and force of movement made by body parts.

Interoceptive functions are located within the body, and concern sensing states of internal organs. Interoceptive functions will tell a person how he feels regarding health.

2.1 Vibrational mechanoreception

General mechanoreception consists of the capability of texture perception, perception of form and the perception of vibration. When conducting research involving vibrational mechanoreception one must know that the spatial and temporal resolutions of reception are not the same throughout the entire body. The spatial acuity can be measured using the *two-point limen* which holds the minimal separation between two stimulated points to be perceived differently.

In the perception of vibrations, two types of stimuli can be discriminated. Stimuli having a frequency below 40 Hz are commonly referred to as flutter. Stimuli with a frequency higher than 40 Hz are denoted as vibration.

These two vibrotactile stimuli are perceived by humans with four types of cutaneous afferents. Two of which are slowly adapting, SA-I and SA-II, and two rapidly adapting, Pacinian Corpuscle (PC) and Ruffini Afferent (RA). The RA is involved with flutter, whereas the SA-I and SA-II are involved with higher frequency stimulation.

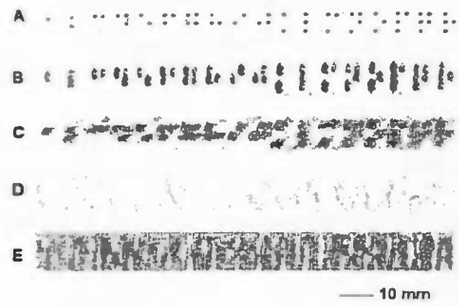

Figure 2.2: Spiking patterns of peripheral, afferent neurons (Zigmond et al., 1999)

Figure 2.2 shows a good overview of the behaviour of each type of peripheral neuron. In this figure, A represents the input signal, a line of Braille. SA-I and RA follow closely the input pattern (B, C respectively), whereas SA-II and PC respond to different properties of the stimulus, in such a way that it causes noise on the input signal (D, E respectively).

2.1.1 Neural coding of vibrotactile stimuli

Talbot (1968) was the first one to discover the way neural encoding in flutter is performed. Talbot found out that the RA and PC respond periodically at the frequency of the stimulus. He also discovered that the RA's hardly change in firing rate when the frequency of a stimulus increases. From both Talbot and Mountcastle we can conclude that the RA and PC systems encode temporal information.

2.1.2 The somatosensory cortex

All the afferents carry their information to the somatosensory cortex, which consists of two main areas, S-I and S-II. Although we know from Mountcastle et al. and Talbot the working of afferent neurons, it is not clear how the somatosensory cortex represents vibrotactile stimuli. Regarding this working there are three possibilities of neural representation:

Spatial encoding. Like the topological map of the primary visual area in visual orientation stimuli (see Kamitani and Tong, 2005; Haynes and Rees, 2005). Since the flutter and vibration are processed in different areas, this could also be the case for subdivisions of flutter or vibration on a smaller scale.

Periodicity encoding. It could also be true that bursts of neural spikes occur in synchrony with the pulses of the vibration.

Rate encoding. The number of spikes in a certain time interval could define the vibrational frequency.

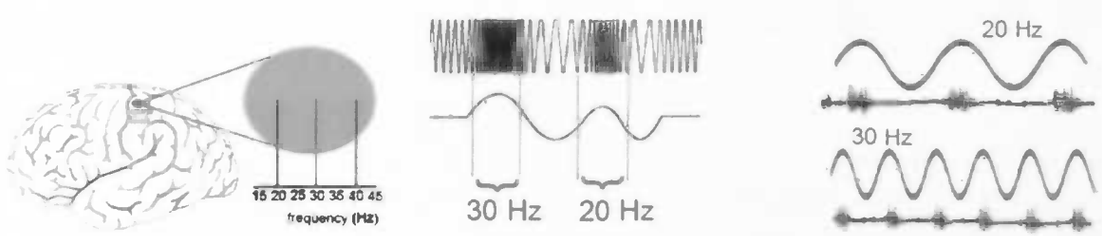

(a) Spatial encoded representation (b) Rate encoded representation (c) Periodicity encoded representation

Figure 2.3: Three possible cortical representations of vibrotactile stimuli.

A major research effort has been done by Mountcastle et al. They managed to correspond the neurometric responses of monkeys with their psychometric responses. By training these monkey to perform the frequency discrimination task, they showed that monkeys perform as well as humans do. They also made conclusions regarding the encoding of postcentral quickly adapting neurons. Mountcastle et al. think that the basis for discrimination is periodicity. After harmonic analysis of the neural spiking the found frequency corresponds exactly with the stimulus frequency.

On the other hand, latest research into the somatosensory cortex of primates have investigated this neural representation further and concluded that rate encoding is significantly correlated with the performance of frequency discrimination. For flutter stimuli the periodicity component is extremely high in the S1 areas but show little effect on this performance (Salinas et al., 2000).

Romo et al. made similar conclusions in his research. Namely, that rate encoding is more likely to be the basis of frequency discrimination. Their results imply that spike periodicity does not play a functional role in the performance (Hernandez et al., 2000; Romo et al., 2003). In the end it is thus still not clear how vibration is represented in the somatosensory cortex.

2.2 Cortical representation of vibration

During this project I have investigated how vibrational stimuli are represented in the brain. Since there are three possibilities of this neural encoding (spatial, frequency or rate), Attempts have been made to find a spatial topology within the somatosensory cortex. This would infer that frequency is represented spatially. If this is not the case, this leaves frequency or rate encoding to be possible representations, eliminating the spatial representation.

A full detailed description of this experiment will be explained in chapter 6. In order to find this spatial topology we will need advanced methods to classify brain states. Support-Vector Machines are currently one of the best method to analyse spatial patterns. For this reason I have used Support-Vector Machines to look at data which has been collected by a Magnetic Resonance scanner. The subsequent chapters provide a brief overview about the principles of fMRI, followed by a description about Support-Vector Machines.

Chapter 3

Magnetic Resonance Imaging

Before going in depth with the research of analysing image sequences produced by a Magnetic Resonance scanner, it is important to have knowledge about the underlying methods and principles in the field of Magnetic Resonance Imaging (MRI). The following section will give a brief overview on these methods and principles, going from general MRI techniques towards the specific kind of MRI used in this project, functional MRI.

3.1 Principles of Magnetic Resonance Imaging

Magnetic Resonance can be best described by measuring the interaction between the nucleus of an atom and an applied external magnetic field. This interaction depends on the environment of the nucleus and therefore has different properties in different tissues.

The interaction only works for polar nuclei, meaning they have a non-zero angular momentum. This is only the case for nuclei with an odd number of nucleons (protons and neutrons). Within medical imaging, we usually look at the signal from the hydrogen nucleus attached to the water molecule, since the human body consists mostly out of water.

There are several steps involved in the acquisition of a magnetic resonance image, which are described below. Firstly, commencing a scan requires the alignment of every nucleus along an axis to a defined plane. This alignment is done by applying a static magnetic field, called the B_0 field. This aligns the magnetic moment of the nucleus either parallel or anti-parallel to the B_0 field, corresponding to two energy states. Since energy is proportional to frequency for electromagnetic radiation, each energy state has an associated frequency.

The difference in frequency between the two states is described in the Larmor equation:

$$v = \gamma B_0 \tag{3.1}$$

where v is the frequency in MHz, γ is the gyromagnetic ratio in MHz/Tesla for the spin under consideration. This equation also relates the frequency of the resulting signal to the B_0 field (Jezzard et al., 2002).

All MRI experiments comprise the perturbance of the nuclear magnetic moment into the transverse

plane using a brief radiofrequency pulse (the B_1 field) and then to measure the signal in a receiver coil. What will be observed is an oscillating signal that decays under an exponential envelope (Jezzard et al., 2002). This decay is termed the *free induction decay (FID)*. For relaxation of the protons, a delay is required in between the observations. This delay is referred to as the *scan repeat time (TR)*.

3.2 T2* relaxation

. Spin dephasing in the transverse plane is influenced by local field inhomogeneities. These inhomogeneities alter the frequency of the proton's precession and speed up the relaxation process. Inhomogeneities depend upon the local physiological state of the brain in particular the composition of local blood supply, meaning the T2* is an indirect measurement of neural activity. (Logotethis and Wandell, 2004)

3.3 The BOLD Contrast Mechanism

The measured composition of the blood flow using the T2* parameter is caused by a physiological process which is called the blood oxygenation. Neurons require oxygen, which is delivered through the haemoglobin in our blood. Oxygenation is the process which takes oxygen from haemoglobin to feed the neurons. After the consumption of oxygen, the haemoglobin changes into deoxyhaemoglobin, which is paramagnetic and therefore alters the MRI signal.

The changing level of oxygen composition is referred to as the *Blood Oxygenation Level Dependent (BOLD) response*. This principle has been known since 1990, when Ogawa et al. discovered that the contrast of a blood vessel varies with changes in the blood oxygenation level.

3.3.1 Neural activity

During the activation of brain cells there are three important parameters involved which cause for variation in the detected MRI signal. When *oxygen metabolism* increases, the *blood flow* increases, which causes a spreading of the deoxyhaemoglobin. With this increase of flow, the vessel expands, which gives a larger *blood volume*, which compensates for the decrease of deoxyhaemoglobin by spreading. Because blood flow is the dominant effect, a net increase of deoxyhaemoglobin will occur during oxygen metabolism. This effect is a delayed version of the neural activity and can be modelled by convolving the activity with a haemodynamic response function (HRF).

Chapter 4

fMRI Analysis

4.1 Introduction

Beside knowledge about the principles of MRI, knowledge on analysing the captured data is required. fMRI analysis enables us to make inferences on the working of the brain.

For the last decade, functional MRI has been a growing field. Whereas researchers started in the early years with validation studies on this technique using lower level functions, the field has now moved on to also endeavour to understand higher cognitive functions.

4.2 Univariate Analysis

With the growth of functional MRI, reliable statistical tests have been designed to make inferences from. These tests can be divided into two categories:

Parametric comprises an approach using assumptions about the underlying distributions in the experiment.

Non-parametric requiring minimal assumptions for validity mainly based on permutation testing.

4.2.1 Parametric designs

This type of design makes use of several parameters because regional physiology will vary systematically with the degree of cognitive or sensorimotor processing or deficits thereof (Friston, 2005). This method of analysis is widely known as the *General Linear Model (GLM)*, which can be stated in the following equation,

$$y = X\beta + \epsilon \tag{4.1}$$

In this model the observed response y can be explained using the matrix X of explanatory variables and a well-behaved (normal) error term ϵ (Friston, 2005). The GLM is therefore a univariate method, since it does not take spatial relations between voxels into account. It is merely a student's t-test calculated over each voxel given a certain time course. These values can then be plotted onto the

structural image of the brain with a certain threshold for significance, so called *p-value*. This p-value should generally be lower than 0.05 to be significant and should be corrected for multiple comparisons, since there are several thousand voxels in the brain over which the t-test is calculated.

The obtained image depicts the overall event related activity for this linear model.

4.2.2 Non-parametric designs

The shortcoming of parametric designs coming from the assumptions on distributions of functional MRI data gave rise for non-parametric approaches (see Holmes, 1994). The most commonly used non-parametric approach is permutation testing. This technique comprises a formal method quantifying the event related activation for a voxel similar to the parametric design, though without making use of any distribution or assumptions based hereon. Within this method, for every combination of all predefined active and rest states the probability for each configuration is calculated (Nichols and Holmes, 2001). The configuration with the highest likelihood will then represent the most plausible model.

4.3 Multivariate Analysis

Since both parametric and non-parametric designs do not take spatial correlations into account, a relatively new approach has emerged using multivariate statistics. This method is more suitable when the intensity of the BOLD signal stays the same over two different states, but certain distinguishing patterns of activation are likely to occur. These patterns could be coactivation or representing neural topology.

Fisher's Discriminant Analysis (FDA) and *Canonical Variates Analysis (CVA)* are commonly used in multivariate approaches. These methods involve rank testing of matrices, checking the linearly independence of columns. Support-Vector Machines (SVMs) have only been introduced since the beginning of 2000 as new approaches in functional MRI analysis. More about SVMs and its application in fMRI research follows in chapter 5.

Chapter 5

Support-Vector Machines in fMRI research

5.1 Introduction

In order to classify brain states, Support-Vector Machines have been used to analyse patterns of activity in the somatosensory area. Since the General Linear Model is not able to detect spatial relations (see chapter 4), this method of analysis will not be able to give information about whether vibration is spatially represented. Although the GLM is able to provide information if the variations were on the order of at least a voxel, the variations in the somatosensory area will lie within a voxel.

This chapter provides a description and past applications of Support-Vector Machines, starting with its history.

5.2 History

Throughout the history of scientific research, quantitative statistics have played an eminent role for the purpose of finding empirical evidence. In cognitive science, researchers have made attempts to solve learning problems with the development of statistical procedures. In this development of research, four periods can be defined (Vapnik, 1995) :

1. Creating the first learning machines
2. Defining the underlying theories
3. Constructing neural networks
4. Finding alternatives for neural networks

In the early years of learning theory, McCulloch and Pitt (1943) and Hebb (1949) have formulated a mathematical description for the binary neuron function, reflecting biological neurons:

$$y = \phi\Big(\sum_{i=1}^{n} w_i x_i + b\Big) \tag{5.1}$$

Here x_i represents the input vector, y_i the output vector and w_i the learning weight vector. At every iteration this weight vector is adjusted according to the difference of the target and the actual y value.

This perceptron model was a major accomplishment in Computer Science, since this was the first system which had the capability to learn unsupervised. After this development, a vast growth of the field of learning theory has lead to a variety of learning systems. Along with these new systems, researchers also founded corresponding theoretical frameworks. The theory of Solving Ill-posed problems and the Empirical Risk Minimisation became the basis for any well-constructed learning machine.

5.2.1 Theory of Solving Ill-posed problems

With any learning problem one assumes that a unique and optimal solution exists and that this solution depends continuously on the available data. These constraints, stated by Hadamard in 1902, form a well-posed problem. In the case where one constraint is lacking, we are dealing with an ill-posed problem. Most of the problems in learning theory are ill-posed and are inverse problems having multiple solutions. Regularisation techniques have therefore been constructed. Regularisation can be seen as a trade-off parameter between fitting the available data or complying with the well-posed aspects of the problem. This is usually done by making assumptions about the smoothness of the solution function.

5.2.2 Empirical Risk Minimisation

Given a two-class problem, the performance of the classifier with decision function $f(x)$ can be calculated accordingly.

In finding the optimal function y for a given input x, the loss of this function is defined as $L(y, f(x, \alpha))$. The function of risk can then be defined as:

$$R(\alpha) = \int L dF(x, y) \tag{5.2}$$

The goal is to find the function $f(x, \alpha)$ which minimizes $R(\alpha)$. Furthermore the original statistics of the distribution $F(x, y)$ is unknown and are reflected poorly in the training set (Vapnik, 1995).

Because of these drawbacks, a specific version of this expected risk has been developed, making only use of the available data. This specified version based on empirical evidence is called Empirical Risk Minimisation:

$$R_e(\alpha) = \frac{1}{l} \sum_{i=1}^{l} Q(z_i, \alpha) \tag{5.3}$$

Here $F(x, y)$ is replaced by $Q(z_i, \alpha)$ and adapted in a way that it now calculates the specific loss function Q for the given sample pair z_i consisting of x and y. Over all the given samples pairs, this gives R_e, the empirical risk.

This description is very important in learning theory because it assumes that in finding the specific minimisation of the training sample, the more general minimisation is found as well. In this assumption

the ERM incorporates the Maximum Likelihood method, trying to find parameter α.

Maximum Likelihood method

The Maximum Likelihood method is used to make inferences about parameters of the underlying probability distribution of a given dataset. For a given probability distribution D, the likelihood function can be defined as $l(\Theta) = f_D(x_i|\Theta)$, providing information about the likelihood of the data given a model Θ. The maximal value of l for Θ, $\hat{\Theta}$ is defined as the Maximum Likelihood Estimator (MLE).

This MLE has desirable properties in statistics. It can be considered to have no bias and to have approximate normal distributions.

5.2.3 Revival of Neural networks

In 1986, the backpropagation was developed by Rumelhart et al. This was a major improvement of the perceptron model. The simple sign function is replaced by the sigmoid approximation. This sigmoid now is a gradient term with respect to all coefficients of all neurons. Using the evaluation of this gradient, gradient-based error minimisation algorithms, such as hill-climbing, can be used in finding minima. With this approach to solve ill-posed problems, overfitting became the next problem. Overfitting stands for the bias towards a learned training set, which deteriorates its performance in general. In the 1990's more attention again was focused to alternatives of neural networks, such as the Radial Basis Functions methods and especially Support-Vector Machines.

5.3 Support-Vector Machines

The mathematical construct of a support-vector machine rises from Vapnik (1995). The main principle for Support-Vector Machines is finding the optimal decision boundary for all dimensions which separates two classes best. This decision boundary, commonly referred to as hyperplane, is optimal when it has maximal margin between the data points and is therefore unique (Vapnik, 1995). Its underlying framework also allows various other learning machines to be constructed.

Consider a two-class classification problem. The classifier has to find the function $f(x)$ such that $y = f(x)$. The performance of this classifier can than be expressed as:

$$E(x, f(x)) = \begin{cases} 1 & \text{if } y = f(x) \\ 0 & \text{otherwise} \end{cases} \tag{5.4}$$

Using Empirical Risk Minimisation (see section 5.1.2) and Structural Risk Minimisation, the Support-Vector Machine tries to find the best mapping using α. The Structural Risk Minimisation defines the trade-off between the quality and the complexity of the approximated function (Vapnik, 1995).

5.3.1 Optimal hyperplane

The optimal hyperplane is the plane $(w \cdot x_i + b = 0)$ which separates the two classes best. The hyperplane needs to have the following mathematical properties:

$$(w \cdot x_i) + b \geq \quad 1 \quad \text{if } y_i = 1 \tag{5.5}$$

$$(w \cdot x_i) + b \leq \quad -1 \quad \text{if } y_i = -1 \tag{5.6}$$

$$\tag{5.7}$$

Given the two classes the function for the optimal hyperplane can be constructed as follows:

$$\phi(\omega) = \| w \| \tag{5.8}$$

In finding the optimal hyperplane, we minimise 5.8 by making use of Lagrangian theory:

$$\pounds = \frac{1}{2} \| w \|^2 - \sum_{i=1}^{l} \lambda_i([(x \cdot w) + b]y_i - 1) \tag{5.9}$$

where λ_i are the Lagrange multipliers. Minimising $\pounds$ to b and ω, means finding the optimal value ω_0 for ω, λ^0 for λ and b_0 for b. This can also be written as the following three equations from which one obtains the properties of the optimal hyperplane:

$$\sum_{i=1}^{l} \lambda_i y_i = 0 \quad \lambda_i \geq 0 \tag{5.10}$$

$$w_0 = \sum_{i=1}^{l} y_i \lambda_x^0 x_i \quad \lambda_i \geq 0 \tag{5.11}$$

$$w_0 = \sum_{\text{SVs}} y_i \alpha_x^0 k(x_i x) - b_0 \quad \lambda_i \geq 0 \tag{5.12}$$

- The first property holds that all λ's must satisfy the constraints (equation 5.10).
- The second property states that the optimal hyperplane can be defined as a summation of the vectors of the training set (equation 5.11).
- The third property holds that only the training data with nonzero λ terms weigh in this description of the hyperplane and are called the support-vectors (equation 5.12).

The Support-Vector Machine performs a mapping of the input data onto a high-dimensional feauture space Z. After this, the Optimal hyperplane is constructed in Z. This mapping function, $k(x_i, y)$, is called the kernel of the support-vector.

For a linear SVM, the kernel function is a dot product in the input space while the kernel function in a nonlinear SVM effectively projects the samples to a feature space of higher (possibly infinite) dimension via a nonlinear mapping.

5.4 Support-Vector Machines applied to fMRI analysis

5.4.1 Research approach

Integrating Support-Vector Machines in any type of research requires a different approach than with traditional fMRI analysis. The research questions are different when using SVMs compared to standard fMRI analysis.

Where standard fMRI analysis is used for investigating any event-related activities where it is impossible to look at relationships between or within brain areas. This approach is more based on the experimental design and on the neural connections going from the sensory input to the corresponding brain area. Research using SVMs has the ability to look at the patterns of interactivity between or within brain areas. The outcome of the SVM corresponding with this neural behaviour can result in a differentiation of the experimental conditions.

With traditional fMRI analysis the output is a statistical map. This analysis is based on the activation pattern of a single voxel given a certain time course, implying that no relationship between voxels is measured. With the use of SVMs this can be investigated.

The resulting activation map in traditional fMRI gives an overview of the event-related activity per voxel. The learned model of an SVM can give a similar output, although the activation map represents here the level of information on which the SVM discriminates the image classes.

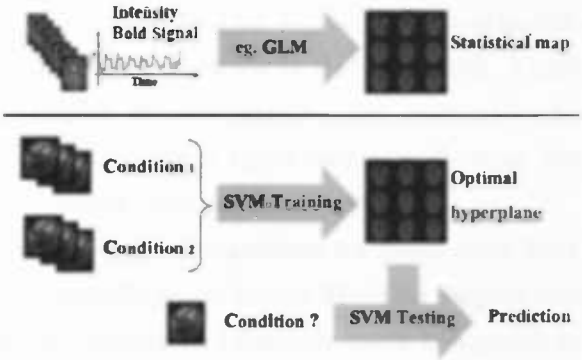

Figure 5.1: Overview of traditional fMRI analysis versus analysis using SVMs

For traditional analysis methods (e.g. Fisher's LDA) the analysis pipeline is continuous from beginning to the end. Whereas with SVMs, the pipeline is two-staged comprising a training and testing phase. Once the training step has taken place, the system will perform testing on a different set of data using the learned parameters from the training step (see Figure 5.4.1).

In the field of functional MRI it has only been since 2003 that support-vector machines are being applied. Despite the short history, interesting projects already have taken place. The following section will give an overview of these projects.

5.4.2 Previous Research

In 2001 Haxby et al. investigated the activation patterns of the entire brain in the experiment showingdifferent categories of objects as stimuli. All presented items belonging to a certain category could be discriminated by looking at the mean within-category and between-category correlations for active neuron populations (see Figure 5.4.2). A population encoding based on the patterns of responses has the capacity to produce a large variety of unique representations of virtually unlimited object categories. Their findings of consistent topographic arrangement may provide a key for decoding the information that underlie face and object recognition (Haxby et al., 2001).

Figure 5.2: Mean within-category and between-category corrrelations three of the eight categories. These correlations are calculated for event-related activated populations of neurons across six subjects. (Haxby et al., 2001)

An imitation of this study, performed by Cox and Savoy, investigated support-vector machines by setting this technique against univariate methods in fMRI analysis. They used the same experimental set-up as in the study from Haxby et al., instead using support-vector machines to classify each category. With this method they showed a significant classification improvement compared to the Linear Discriminant Analysis method Cox and Savoy (2003).

Mitchell et al. explored several types of classifiers on fMRI data from different studies. With this exploration he showed that the linear Support-Vector Machine outperforms all of the nearest neighbour algorithms and that the accuracy increases relatively more compared to the Gaussian Naive Bayes classifier when the dimensions of the data reduces.

In 2005 there has been a major breakthrough in the neuroimaging community because Kamitani and Tong as well as Haynes and Rees convinced researchers on a wider scale that the use of multivariate techniques are easy to use and for certain group of problems well capable in providing answers.

Kamitani and Tong investigated the distribution of the visual cortex when applying different orientation stimuli. They showed that the SVM was able to find a spatial preference map of V1 and that the shown orientation could be read from the fMRI signal (see Figure 5.3).

Haynes and Rees used Fisher's Linear Discriminant Analysis to predict the presented visual stimulus orientation from the primary visual cortex. They even were able to predict this where the subject was unaware of having seen the stimulus.

Strother et al. designed a data analysis framework for the assessment of prediction accuracy versus

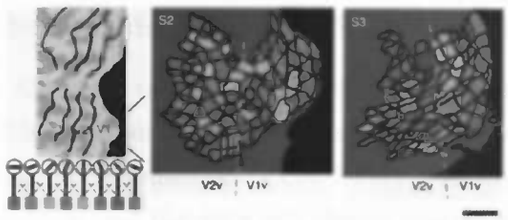

Figure 5.3: Orientation preference of V1 and V2 (Kamitani and Tong, 2005)

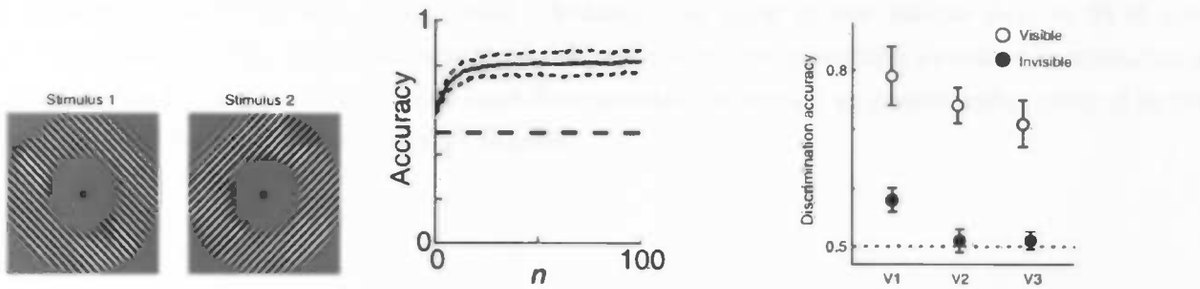

(a) Visual stimuli (b) Prediction accuracy (c) Prediction of unaware stimuli in V1

Figure 5.4: Results from Haynes and Rees (2005)

signal-to-noise ratios. This framework consists of resampling and cross validation methods and can be used as a validation tool for testing and optimizing methodological choices in functional neuroimaging (Strother et al., 2002). In 2005, LaConte et al. reported on the use of SVM in temporal block design with regard of this framework. They showed, by making use of the validation tools from Strother et al., that a SVM with a linear kernel outperforms the SVM with a n-polynomial kernel.

(a) Visual stimuli (b) Orientation

Figure 5.5: Results from Boynton (2005)

From the review of Boynton, Figure 5.5 shows that a presented generated pattern (5.5(a)) in their experiment results in selectivity inside each voxel (5.5(b)).

Mourão-Miranda et al. also applied SVMs to fMRI analysis. Their research focused on whole-brain volumes from an attention study. The experiment conducted was a face matching task and a location matching task. Firstly, they also showed that SVM outperforms LDA in robustness and performance. Secondly, they were able to classify, with an error rate of 20 percent, brain volumes corresponding to the face-matching task from the brain volumes corresponding with the location-matching task.

5.4.3 Different approaches within SVM based fMRI research

Within the previous research, three approaches can be extracted on the method of applying SVMs. Namely, the application of SVMs based on either raw data or data preprocessed by feature selection or data preprocessed using dimensionality reduction (PCA/SVD). The approach where feature selection has been applied, has the limitations that the data will always contain a bias because feature selection is done by humans or machine learning algorithms, based on a priori hypotheses. This in contrast with the application of dimensionality reduction, which has the advantage that it does not use a priori information (Mourão-Miranda et al., 2005). Although the input of raw volume data in SVM does not suffer from any bias, the disadvantage of this method is that it requires extensive computational resources. For these reasons, we mainly used dimensionality reduction as preprocessing method in this study, and will be explained in the next chapter.

Chapter 6

Methods

6.1 Experimental design

The goal of this project is to predict the vibration frequency presented to the subject performing a vibrotactile discrimination task using Support-Vector Machines. Therefore, a two-alternative forced-choice task has been designed which had to be performed inside a scanner (Green and Swets, 1966).

This task consists of classifying a range of frequencies as being lower or higher than a certain reference frequency presented in advance of this stimulus. The reference stimulus has been set to 30 Hz, and the frequencies which had to be classified in respect to this were 27, 28, 29, 29.5, 30.5, 31, 32 and 33 Hz. A stimulus presentation consists of two seconds of reference stimulus followed by four seconds of a frequency in the range mentioned above. In order to get the BOLD level to a resting state and for the ease of the subject, a resting period of fourteen seconds was inserted after each stimulus.

In this experiment we have defined the following terms:

Trial consisting of the reference frequency followed by the stimulus frequency and concluded with the resting period. The length of each trial will then last for 20 seconds.

Run comprise the trials containing the eight stimulus frequencies randomly distributed so every stimulus is presented five times. The time a run takes is approximately 13 minutes.

Session contain four runs each having a different distribution of the trials.

Data was acquired from three sessions of two subjects. For a third subject, data from only one session was acquired, after which the decision was made not to analyse, since results of this session had large artefacts in the images due to subject movement and errors in the psychometric results which could not be corrected for.

6.2 Data Acquisition

6.2.1 Image acquisition

Image data was acquired using a 3 Tesla Siemens scanner. A two-dimensional axial echo planar imaging sequence was used. This results into the acquisition of 28 slices through the brain every 2 s.

The matrix size was 64x64 pixels and the field of view 22.4 cm, giving an in-plane voxel size of 3.5 mm.

6.2.2 Psychometric acquisition

The responses of the subject were captured using a response box connected to a computer. This response box contains four buttons, two of which were unused. The possible answers were *lower* or *higher* (see also Figure 6.1(b)).

6.2.3 Stimulus presentation

For the vibrational stimulus we used a *tactor* connected to a pulse generator, which in its turn is controlled by a program written in Labview[1]. This program allows us to set the gain of the pulse as well as the frequency. this pulse is converted to a flutter inside the tactor.

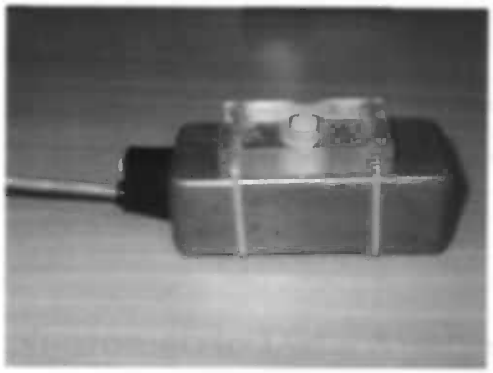

(a) The tactor

(b) Example setup

Figure 6.1: Equipment used in the experiment

Before the tactor can be used in the experiment in a valid way, the tactor requires subject-specific calibration. Within this calibration the least detectable vibration is found by homing in on this vibration by varying gain For the stimulus presentation the gain was set to three times the detection threshold.

For the presentation of screen instructions we used Presentation[2]. This program was also connected to the response box inside the scanner registering the button presses of the subject. This response box contains four buttons of which two were used to store the answer, being either lower or higher.

[1] See National Instruments, http://www.ni.com/labview
[2] See Neurobehavioral Systems, http://www.neuro-bs.com/nbs_online/presentation/

6.3 Psychometric Data Analysis

After the acquisition of the psychometric data, performance rates were calculated and modeled using Matlab[3]. To analyse the performance of subjects in this experimental task, Receiver-Operating-Characteristic (ROC) graphs have been drawn. ROC graphs, frequently used in signal detection theory in psychophysics. An empirical ROC graph captures the performance of a subject over a certain range of ordered presented stimuli over a certain number of trials. Thus the curve in the ROC graph may change if the physical conditions of the experiment varied (Green and Swets, 1966). After applying the central limit theorem which states that the sum of a large number of independent variables approaches a normal distribution the underlying distribution can be perceived as normal, since we have a large number of independent variables.

In finding this ROC function, the nonlinear least-squares fitting has been used. The function of the ROC in this model holds:

$$f(x) = 50\ \mathrm{erf}(\alpha x) + 50 \tag{6.1}$$

Where α defines the steepness of the curve, which represents the performance of the vibrotactile task. The steeper the curve, the better the classification was performed. The error function erf is defined as:

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt \tag{6.2}$$

6.4 Neurometric Data Analysis

6.4.1 Data Preprocessing

For standard fMRI preprocessing, the acquired volumes were imported into BrainVoyager[4]. This program (written by prof. dr. Rainer Goebel at the University of Maastricht) is generally being used to perform analysis and visualization in fMRI studies. Standard preprocessing consists of four components:

Linear Trend Removal This preprocessing step removes any slow trends caused by the heating of the gradient coils during the scan.

3D Motion Correction This feature corrects for high frequency motion in any direction caused by subject movement.

Slice scan time correction correcting for the fact that slices of one image are acquired at different points in time. The nature of this error stems from the technique used to create a 3D image inside a scanner.

Temporal filtering performs a high pass filter in each timeseries of a voxel to remove any low-frequency physiological fluctuations, such as changes in blood pressure.

[3] See Mathworks, http://www.mathworks.com
[4] See Brain Innovation BV, http://www.brainvoyager.com

After preprocessing, the fMRI images are mapped onto a structural image. This is required since, for each session, the fMRI images will have a slightly different location in the head of the subject. Therefore mapping it to the high-resolution structural volume with specific alignment parameters (rotational and translational) results in having all fMRI sequences in the same space. Finally, the fMRI images projected in the space of the structural volume are exported to the Analyze format. This file format is the default format in the fMRI community and is therefore a portable format.

6.4.2 Temporal Compression

The sequence of images can be temporal compressed according to the experimental paradigm. We chose to average three images (i.e. 6 s of data) corresponding to the time one vibrotactile stimulation was applied. Additionally, The delay and diffusion of the haemodynamic response was taken into account here, so the actual temporal boundary for every condition was shifted forward by two volumes (i.e. 4 s). After this compression, the average image from the previous resting period was subtracted from each condition. This results in an image for every condition according to our experimental design.

6.4.3 Software Framework

For this project I have made use of the multiSVM package written by Joachims. It has been extended and adapted in such way that fMRI images can be dealt with. Furthermore, the Gnu Scientific Library (GSL) has been used to embed Singular Value Decomposition.

6.4.4 Singular Value Decomposition

Since the number of voxels is larger than the number of scans, the number of degrees of freedom is also large. *Singular Value Decomposition (SVD)* allows us to reduce the dimensionality of the observation space and hereby lower the number of degrees of freedom(Kjems et al., 2002).

The dimensions of the dataset in our case are the number of voxels (114688) by the number of volumes, which only varies from 40 to 1800. SVD corrects for this overdetermined matrix and returns a compressed approximation of the original dataset.

The compressed dataset which SVD returns is more robust to numerical error. SVD converts a given dataset X into three matrices, S, V^T and U provided equation (6.3). X has been corrected for, by subtracting the precomputed mean per voxel over time.

$$X = USV^T \tag{6.3}$$

This decomposition method extracts the eigenvalues of X and stores them in S. U and V^T are called left and right singular vectors, respectively. Depending on the dimensions of X the projected dataset X_P can be obtained by multiplying X with the left singular vector V for X having $M > N$, and with the right singular vector U for X having $M <= N$. This X_P will have the size of $M \times M$.

This method is used to reduce our dataset without losing statistical properties of the dataset, by importing X_P into the Support-Vector Machine instead of X.

To decrease the amount of computation we first created the decomposed matrices with SVD of the complete set, after which train and test sets were generated from the original set and multiplied by U^T, since our number of images is smaller than the number of voxels in all cases. During each run of the SVM classifier, the subset of training and test set are multiplied with U^T to get the projected subset:

$$X_{P_{train}} = X_{train}^T U^T \tag{6.4}$$

$$X_{P_{test}} = X_{test}^T U^T \tag{6.5}$$

These multiplications result in a projected matrix with an item per row. Since the SVM creates the hyperplane in the dimensions of the SVD, it is impossible to see whether the correct features have been picked up by the SVM. Therefore, it is possible to create a backprojected image of this hyperplane, by multiplying the matrix of support-vectors with V^T. The resulting matrix is in the same space as the fMRI images and can be inspected and validated.

6.4.5 Masking Image

To filter out possible other event-related activation of brain regions Two types of masks have been used.

Functional. This mask was based on the activation map created by a GLM model in BrainVoyager. Afterwards a threshold was set so only high activation regions are used in the training and testing of the SVM.

Structural. The second type of mask was based on the location of the somatosensory areas. Because the backprojected image showed that the SVM used the eyes extensively, this had to be ruled out. The image was split into four quadrants, where only the upper right quadrant has been used.

6.4.6 k-Fold Cross Validation

The statistical validation of the classification has been done using k-fold cross validation according to Kjems et al. (2002). Using only one training and testing set, the outcome of this classification cannot be excluded from chance. The entire dataset comprises 20 images per presented frequency per subject per session, which has been split up in consecutive train and test images of a tenfold. These combined subsets give ten sampled sets which have been used for training and testing the Support-Vector Machine. This also results in ten different performances of each discrimination task. These ten scores were averaged afterwards, to obtain an overall performance score.

6.5 Data Analysis

In the analysis of the neurometric results I have used different configurations of the grouping of training and testing data. For each group, the SVM tried to find distinctive properties where the one group can be separated from the other.

In order to assure that the SVM is valid, a validation study has been performed before running the experiments. This study holds the performance of two classes, namely the volumes captured in a period of rest and volumes captured during the stimulus presentation. This is described in chapter 7.

The actual experiments with the SVM can be seen as an extensive simulation of the psychometric experiment. Volumes have been labelled according to the presented stimulus during the session. For each tenfold set 36 volumes per session have been used as training data. For the test set 4 images per session have been used. This results in 40 classifications, of which the average forms the performance of the SVM for the two classes. In accordance with my hypothesis, I expect that the performances based on the classes are related with the frequency difference of the classes. The higher the frequency difference amongst the classes, the better the performance of the SVM. Results of these classifications can be found in chapter 8.

Based on the findings of Cox and Savoy (2003), Mitchell et al. (2004), LaConte et al. (2005) and Mourão-Miranda et al. (2005) only linear kernels have been used with default kernel and default margin settings.

To answer the question whether learning trends can be picked up by the SVM a third type of experiment has been performed. In this method, labelling is based on the time when the volume has been captured. If the SVM performs well on this task, this denotes the presence of a feature in the signal which could represent learning.

Chapter 7

Validation of SVM

7.1 Introduction

Before the SVM can be used to make inferences on how signals are represented in the brain, it requires validation first. This has been performed by testing the SVM on a previous dataset collected at the Magnetic Resonance Imaging and Research Analysis Centre.

7.2 Methods

7.2.1 Validation dataset

The dataset which has been used for validating the SVM holds an experiment regarding thumb movement. In this experiment the subject had to make thumb movements in four variations. Either the left thumb or the the right thumb had to be moved for two different periods of time.

As the results of the General Linear Model show, a distinct spatial map can be made based on which thumb the stimulus was applied.

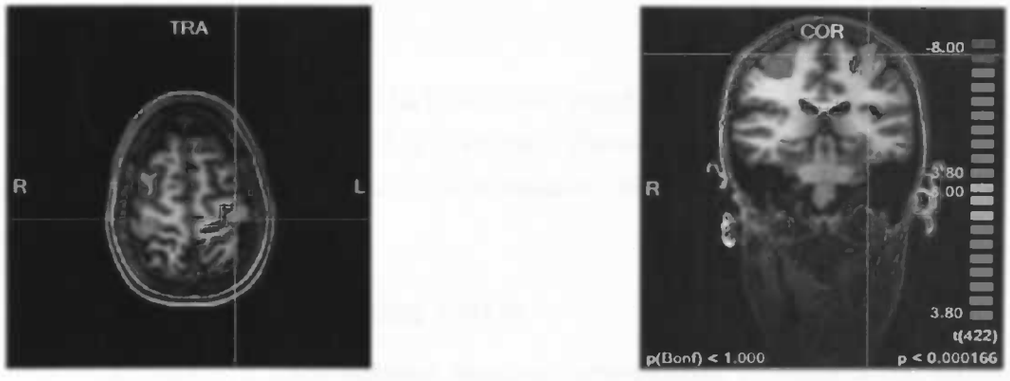

(a) Axial view

(b) Coronal view

Figure 7.1: Analysis of event related activity using the General Linear Model: Spatial dependence of left versus right stimulus presentation. Blue represents the regions where activation caused by left thumb movement was higher than by right thumb movement. Red represents the regions where activation caused by right thumb movement was higher than by left thumb movement.

Whereas the number of stimulations is not spatially represented, it is possible to look at the amplitude dependence.

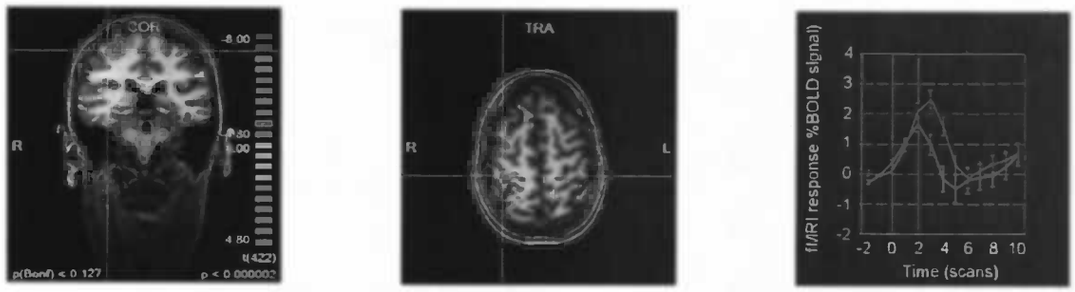

(a) Axial view (b) Coronal view (c) Amplitude graph

Figure 7.2: Event related amplitude curves in the BOLD signal. The somatosensory area shows increased amplitude for five stimuli (brown) in comparison with one stimulus (blue)

7.2.2 Backprojection of the learning vector

A second validation study has been performed to assure the correct working of the software. A backprojected image has been calculated from the most simple classification, where volumes corresponding to stimuli have to be discriminated from volumes corresponding to a resting period. This should classify as hundred percent, because the neural activity patterns are very different for both classes over the entire brain. The calculated learning vector, which is stored in SVD space, can then be backprojected, using the inverse function for SVD, into volume space and mapped onto the structural volume.

7.3 Results

7.3.1 Validation dataset

As Figure 7.3 shows, the Support-Vector Machine is well capable in classifying spatial patterns, but lacks in the ability to discriminate the number of stimuli. The outcome shows that the SVM was able to classify whether the left or the right hand was stimulated, although it shows a weak performance on classifying one versus five puffs of air.

7.3.2 Backprojection of the learning vector

After successful classification of 100% between volumes corresponding to either a rest period or to any stimulus, the backprojected functional volume has been calculated and fitted onto the structural volume. This is plotted in figure 7.4, also showing a plausible setting, where the somatosensory area and several motor areas show up.

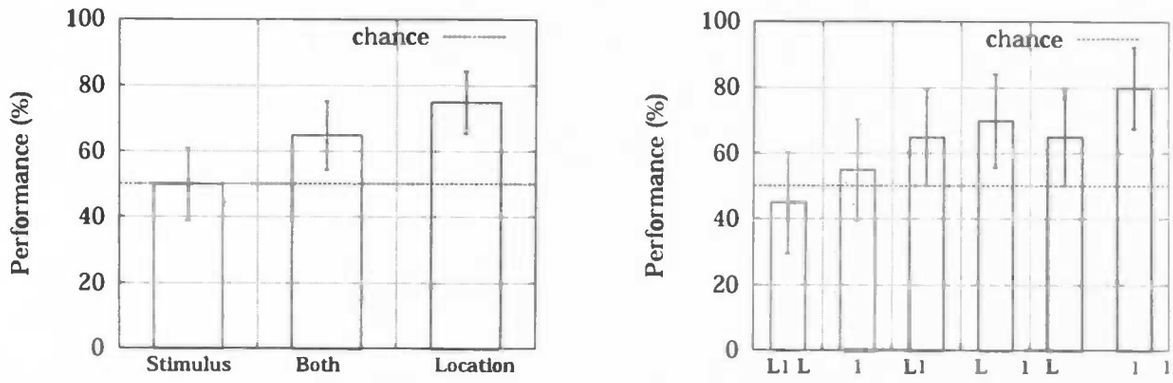

(a) (b)

Figure 7.3: Binary classification results of the validation dataset by the SVM based on whole brain volumes. Both classes differ in either the location of the stimulus (based on 80 experiments), or in the type of stimulus (based on 80 experiments), or in both properties (based on 80 experiments). Figure 7.3(a) shows a categorisation of the types of classification performed in Figure 7.3(b).

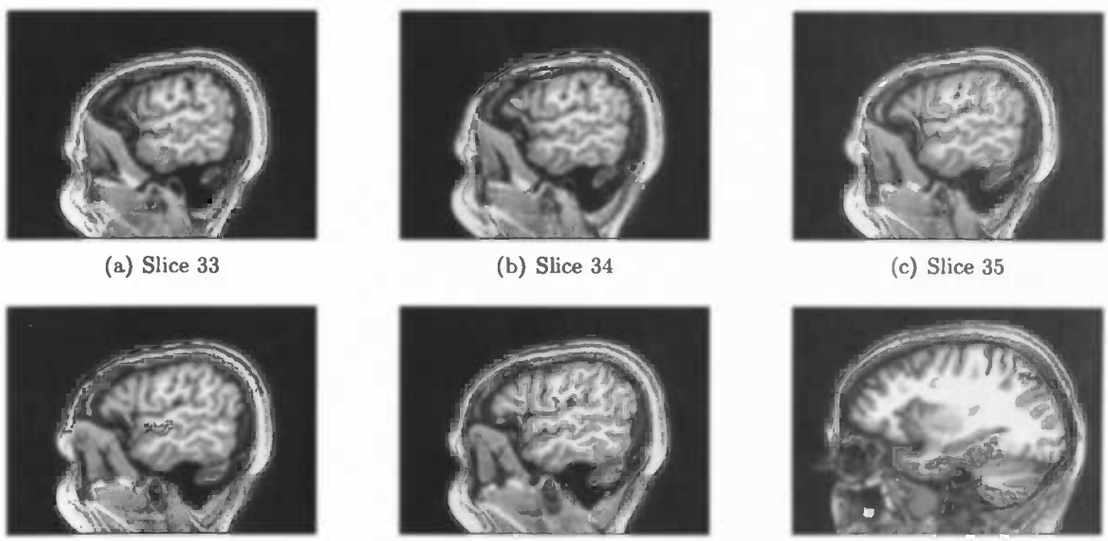

(a) Slice 33 (b) Slice 34 (c) Slice 35

(d) Slice 36 (e) Slice 37 (f) Slice 115

Figure 7.4: Backprojected image of Subject 1. The somatosensory area is picked up as one of the most active parts in the brain.

Chapter 8

Results

8.1 Psychometric Results

The psychometric data has been fitted using ROC graphs, where the steepness of the curve signifies the performance. As figures 8.1 and 8.2 shows, there is a large difference between subject 1 and subject 2. Whereas subject 1 does show an increase in the performance of discriminating the vibrational stimuli over the three scans, subject 2 shows the opposite effect.

8.2 Neurometric Results

Since the psychometric results are not pointing in one direction for both subjects, the neurometric results should show similar trends according to my hypotheses. This would only be the case if a coupling of the psychometrical and neurometrical performance exists and if the discrimination process can be captured using fMRI. The following sections I will present the results for the defined cases as previously discussed in the methods section.

8.2.1 Preprocessing with SVD

Figure 8.2.1 shows the cross-validated average performance of the SVM per session. Each of those ten cross-validation sets consists of 36 training images and 4 testing images. Dimensionality reduction has been applied on each set individually. This dimensionality reduction reduces the input space from the number of voxels to the number of volumes. SVD compresses the original vectors with the size of 114688 elements to vectors having the size of the number of volumes used in the experiment.

8.2.2 Feature selection

Beside preprocessing with SVD, a different approach has been applied. Figure 8.2.2 show the results of classification after feature selection based on an activation map. This activation map has been generated by BrainVoyager using the General Linear Model.

After this preprocessing step, 1057 voxels were used in the SVM. As figure 1 shows, performances of the SVM do not increase overall. The use of the functional mask makes the performance curves

less stable. Still, the SVM is not able to classify above chance which stimulus has been presented. Only from the dataset where Subject 1 performed best, it can be said that with the most difference in frequency the SVM can classify the volumes significantly above chance.

8.2.3 Difference in psychometric performance

The third hypothesis of this study is that learning takes place within the frequency discrimination task. Therefore I have been trying to find trends in both psychometric and neurometric performances over time. Unfortunately, it is not possible to classify scanning session from each other, because of the numerous of parameters which differ. Every session is processed in a slightly different space, which leads to different mappings of functional data onto structural data. Even after concise alignment one cannot make inferences of the results of this classification. However, within a session it is feasible to look at differences over time. I have split up a session into four classes over time, preprocessed the volumes with a functional mask and fed as input to the SVM. The functional mask extracts 1057 most active voxels in the brain. Figure 8.2.3 shows the results of this classification. This figure shows that SVM is able to classify well above chance at which timepoint the volume has been captured.

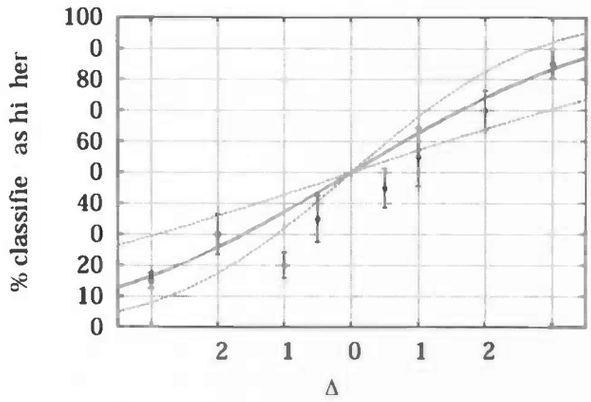

(a) Subject 1: Session 1, $\alpha = 0.229$, $R^2 = 0.84$

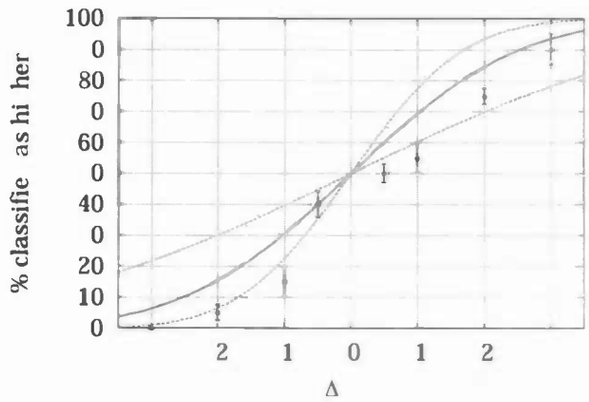

(b) Subject 2: Session 1, $\alpha = 0.359$, $R^2 = 0.89$

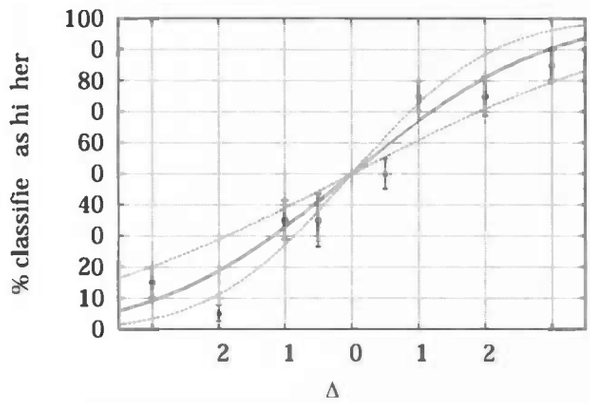

(c) Subject 1: Session 2, $\alpha = 0.313$, $R^2 = 0.92$

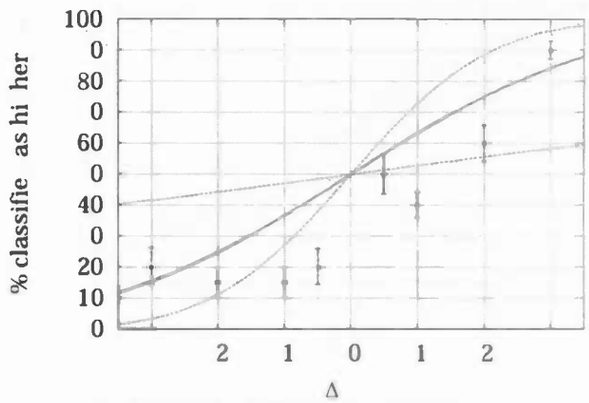

(d) Subject 2: Session 2, $\alpha = 0.24$, $R^2 = 0.52$

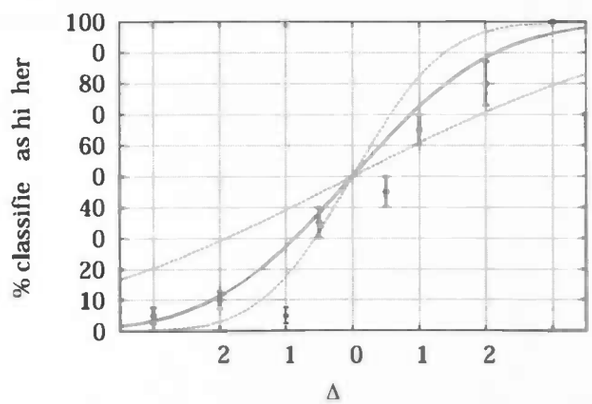

(e) Subject 1: Session 3, $\alpha = 0.426$, $R^2 = 0.89$

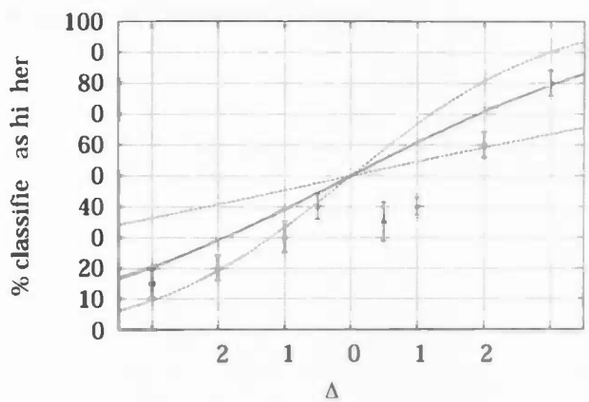

(f) Subject 2: Session 3, $\alpha = 0.195$, $R^2 = 0.70$

Figure 8.1: Psychometric results of Subject 1 and 2, with nonlinear fitted error function, $f(x) = 50\ erf(\alpha x) + 50$ and 95% confidence interval lines. On the x-axis the relative difference of the stimulus frequency with the reference frequency is plotted against the percentage called higher by the subject. These figures give a clear overview that as the stimulus frequency is homing in towards the reference frequency, the task becomes more difficult. Furthermore these graphs show that subject 1 performs better over time, whereas subject 2 shows no increase in performance over the sessions.

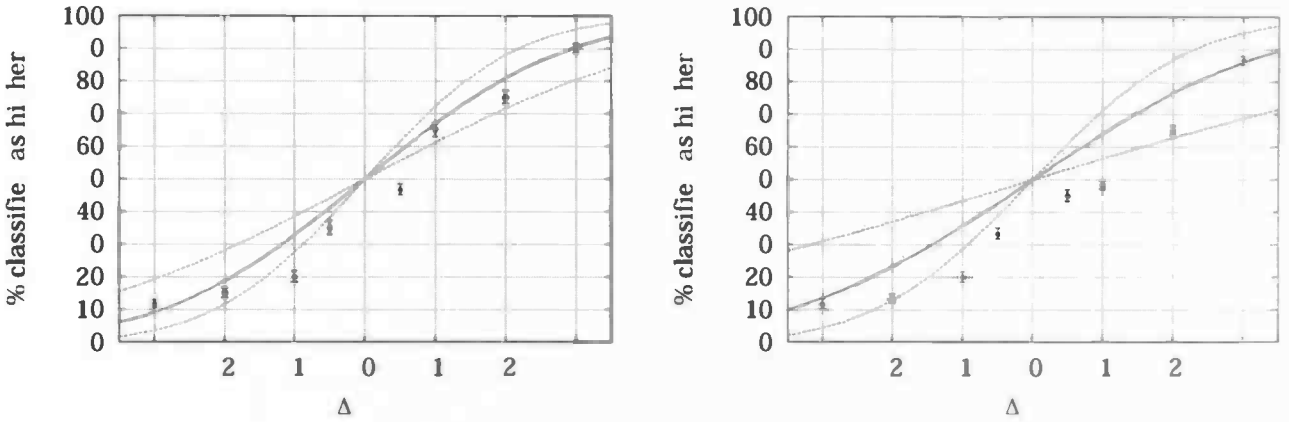

(a) Subject 1: Average, $\alpha = 0.31$, $R^2 = 0.93$

(b) Subject 2: Average, $\alpha = 0.27$,$R^2 = 0.79$

Figure 8.2: Psychometric averages of Subject 1 and 2. These figures show the overall performance of both subjects. On the x-axis the relative difference of the stimulus frequency with the reference frequency is plotted against the percentage called higher by the subject. It is clearly visible that subject 1 performed better than subject 2, since the steepness of the curve is higher.

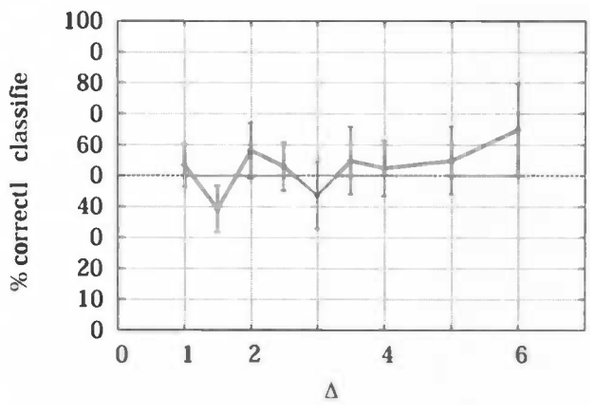

(a) Subject 1: Session 1, preprocessing with SVD

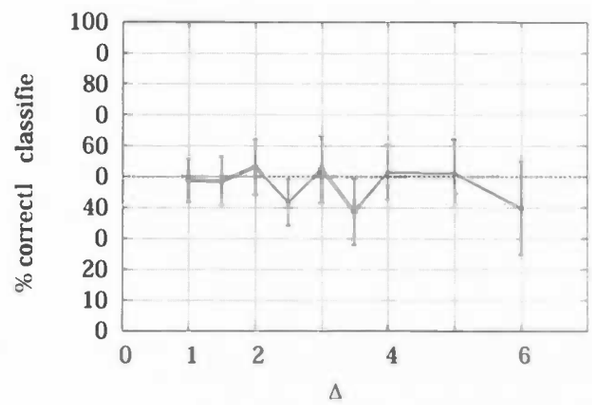

(b) Subject 2: Session 1, preprocessing with SVD

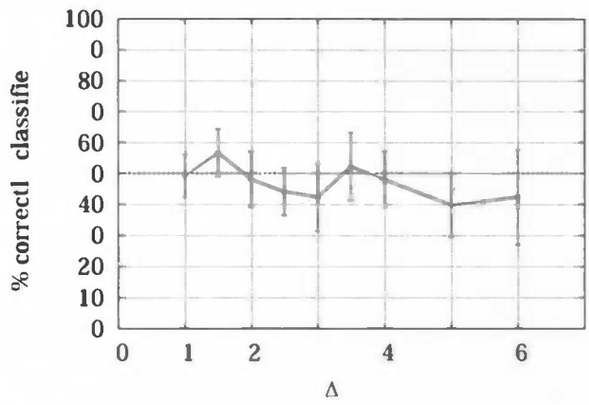

(c) Subject 1: Session 2, preprocessing with SVD

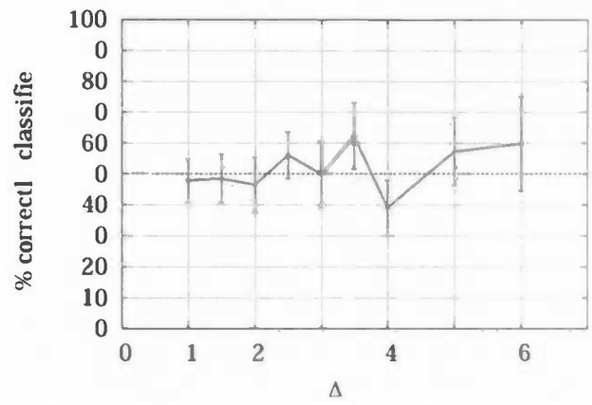

(d) Subject 2: Session 2, preprocessing with SVD

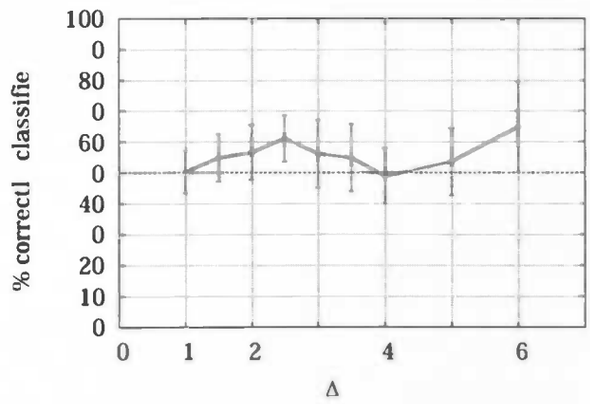

(e) Subject 1: Session 3, preprocessing with SVD

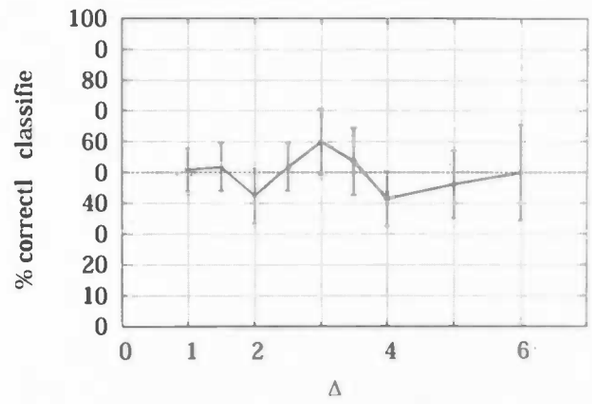

(f) Subject 2: Session 3, preprocessing with SVD

Figure 8.3: These figures show the performance of the SVM on the fMRI volumes preprocessed using Singular Value Decomposition. The volumes have been categorised according to the presented stimulus during the collection (6 Hz difference means classification between 27 and 33 Hz). On the x-axis every possible pairwise classification having the same difference in frequency has been averaged. This has been plotted against the average performance of these classifications. Based on entire volumes, preprocessed using Singular Value Decomposition, no significant difference can be found by the SVM.

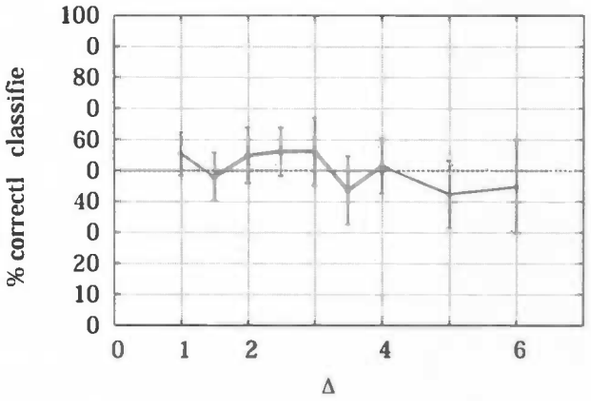

(a) Subject 1: Session 1

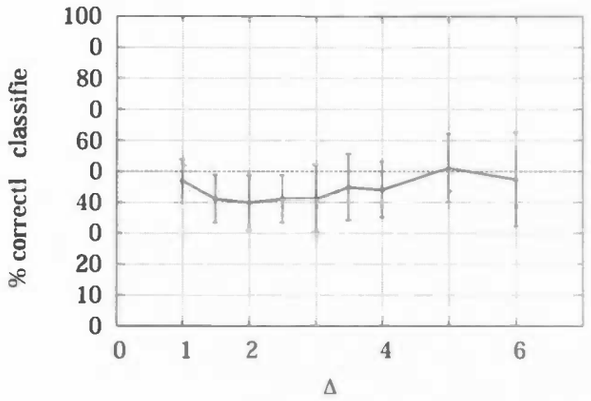

(b) Subject 2: Session 1

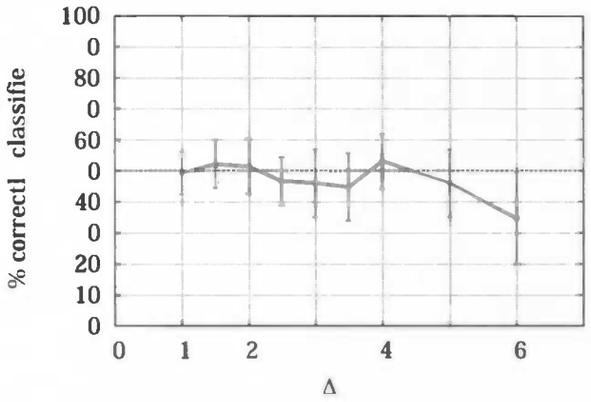

(c) Subject 1: Session 2

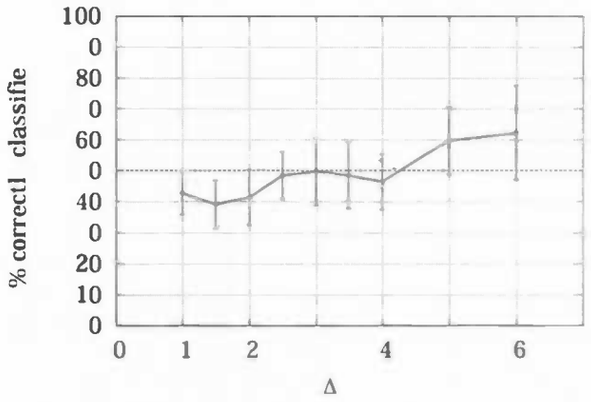

(d) Subject 2: Session 2

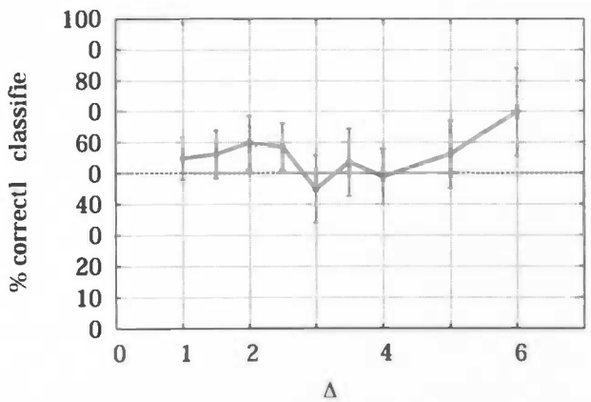

(e) Subject 1: Session 3

Figure 8.4: Performances of the SVM after feature selection based on event related activity map created in BrainVoyager. After this preprocessing step 1057 voxels were taken into account by the SVM. The volumes have been categorised according to the presented stimulus during the collection (6 Hz difference means classification between 27 and 33 Hz). On the x-axis every possible pairwise classification having the same difference in frequency has been averaged. This has been plotted against the average performance of these classifications.

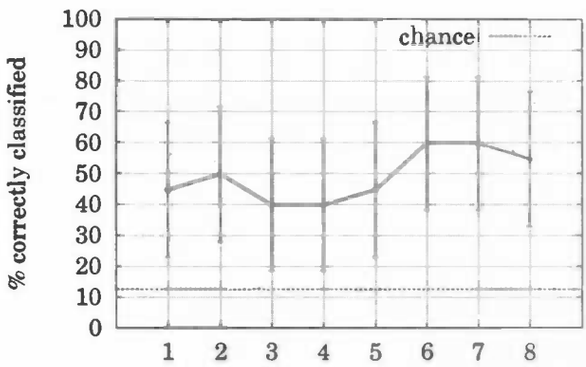

Figure 8.5: Classification performance of SVM on volumes corresponding to eight classes. The classes have been made based on elapsed time from four runs within the third session of Subject 1. The volumes have been preprocessed using a functional mask. The input of the SVM consists of vectors with 1057 voxels.

Chapter 9

Discussion

9.1 General discussion

Although the hypotheses introduced in chapter one have to be withdrawn, this research has confirmed previous research of Mountcastle et al. and Romo et al.. Furthermore, the results should not be seen as a failure of the experiment. The underlying objective of this study was to set-up integrative techniques in pattern recognition for the purpose of fMRI. Now, future projects have the possibility to also use this method beside traditional GLM studies.

The SVM was not able to classify the volumes according to our class definitions. This could stem from either the possibility that the frequency component is not reflected in the BOLD signal, or that the SVM is not able to pick up the signal. In both cases this eliminates the spatial represenation for frequency. This leaves rate encoding or periodicity encoding to be possible representations.

The validation study showed that using GLM, one is able to highlight differences in amplitude of the BOLD signal. Although both in the validation and in the actual experiment the SVM could not pick up this rate encoded difference.

9.2 Specific discussions

This section will give an overview of shortcomings or ideas for improvement for each of the steps during this research.

9.2.1 Experimental design

Number of subjects. More subjects will provide a better overview of the learning trends as well as the relation of the performances of both the subject and the SVM classifier. In this study it is impossible to draw any conclusions from any possible relation, because both subjects showed different learning curves but similar SVM performance.

Range of stimulation frequency. A different range of frequencies could have been used where flutter versus vibration could have been investigated. This would not have given answers to my research questions, because it is already known that the afferents of vibration and flutter

both lead to different brain areas within the somatosensory cortex. If there was a relationship between the location and frequency within flutter, this would have shown in any case.

9.2.2 Psychometric Analysis

Confidence levels. As mentioned in the previous section, a larger pool of subjects will also decrease the error of the psychometric curves. This would still not cover for the individual error over the runs. If a single subject underwent more runs in the scanner this would give a better overview of the individual performance and learning curve.

9.2.3 Neurometric Analysis

Baseline studies. Before starting with SVM an option would be to also perform GLM analysis or even Fisher's LDA to obtain a baseline performance. Although, this would show very little results and would have taken too much time, other methods of analysis could have better shown the strength of SVMs.

Image registration. During the analysis we found that there had been slight shifts in the image registration caused by a problem in Brainvoyager regarding image dimensions. Since the spatial element is crucial for the SVM, reparing this took the necessary time. For future use of SVM, one has to be certain that the preprocessing of the images for every scan are converted into the same dimensions.

Region of Interest. In the beginning of this project the idea was to perform the analyses on the entire brain volume, although at a later stage it appeared that it is better to begin with a small section of the brain (the somatosensory cortex) and then broaden the region of interest to see wether classification can also take place based on whole volumes.

On the one hand the testing of the individual routines of the program starts off with an entire image. Once this works valid, one can go towards programming routines for subselections of this image. On the other hand the outcome of the SVM can only be taken seriously if all other factors for variation have been ruled out. This empowers the researcher to make strong conclusions about the results. The danger lies in the fact that starting with entire brain images, SVMs might give high performances based on different sources than those of interest.

Difference in psychometric performance The discrimination performances of the SVM were particularly high when labelling a scan as a pool. Somehow the SVM finds features which changes over time. Even when splitting a scan into two parts and pooling either part of the scan,the SVM could tell with 70% accuracy from which part the fMRI volume stems from. Needless to say is that those items are left out of the training procedure. In order to investigate whether plasticity might occur over the sessions, more imaging data is required. A solution would be to follow subjects individually over a longer period. In this period, several series of sessions of this task are required to validate learning trends.

9.3 Future research

In the search of finding the representation of frequency in the somatosensory area, further studies can be made using MEG, which is able to capture neural activity at a higher time resolution. If this study would be continued, it must also be noted that human performances on the task have to be high before it can be compared to the performance of the SVM.

The SVM will be used in the investigation of how color is represented in the primary visual area. Since it is known that the Lateral Geniculate Nucleus (LGN) represents primary colours in different spatial locations, the SVM can be validated using datasets of the LGN. The SVM will then be applied to exract mechanisms of color perception in the primary visual cortex, using the prior knowledge of the LGN.

Furthermore, a large variety of applications in brain research can have benefit from this technique.

Chapter 10

Conclusion

At the end of this project we can exclude that the frequency in the somatosensory cortex is spatially distributed based on the results. This leaves two alternatives: either frequency is rate encoded or frequency is encoded by periodicity.

This research has also shown how to integrate recent techniques borrowed from pattern recognition into fMRI research.

The research questions from chapter one will be revisited and conclusions can be drawn based upon the performed research.

1. What is the general classifier performance on large sets of brain images collected either during the application of vibration stimulation or during rest?

   I collected a large set of performances for varying sets of brain images. The SVM was not able to extract the information component regarding the presented stimulus. With the validation study the SVM was able to discriminate rest volumes from active volumes.

2. What is the human performance on the vibrotactile discrimination in comparison to classifier performance on the discrimination of cognitive states?

   In our experiment, Subject 1 showed an overall better performance than subject 2. In the results of the SVM it is not possible to correlate these trends with the SVM performances. The fluctuating performances of the SVM show that the SVM is not able to pickup signals regarding the presented stimulus and therefore do not provide enough stimulus information.

3. What do learning trends of the vibrotactile discrimination task by humans look like?

   I expected to see an increasing trend in human performance, whereas this was only the case for Subject 1. Subject 2 showed a decrease in performance over the sessions. Although it is impossible to capture general learning trends based on two subjects, the results of this experiment shows that there are no general learning trends.

Finally, a conclusive answer can be given to the main question:

Is the support-vector machine able to provide information about the performance of the subject?

The results of the classification of SVM, in order to find a spatial representation, showed that SVM was not able to discriminate volumes based on the presented frequency-stimulus. According to this, it can be stated that there is no spatial information represented in the BOLD signal. This infers that frequency is not spatially represented in the somatosensory cortex.

Bibliography

Boynton, G. M. (2005). Imaging orientation selectivity: decoding conscious perception in v1. *Nature neuroscience*, 8(5).

Cox, D. and Savoy, R. (2003). Functional magnetic resonance imaging "brain reading": detecting and classifying distributed patterns of fmri activity in human visual cortex. *NeuroImage*, 19.

Friston, K. (2005). Brain function modeling. *Annu. Rev. Psychol.*, 56:57 – 87.

Green, D. and Swets, J. (1966). *Signal Detection Theory and Psychophysics.* John Wiley and Sons.

Haxby, J., Gobbini, M., Furey, M., Ishai, A., Schouten, J., and Pietrini, P. (2001). Distributed and overlapping representations of faces and objectes in ventral temporal cortex. *Science*, 293:2425–2430.

Haynes, J. and Rees, G. (2005). Predicting the orienatation of invisible stimuli from activity in human primary visual cortex. *Nature neuroscience*, 8:686 – 691.

Hernandez, A., Zainos, A., and Romo, R. (2000). Neuronal correlates of sensory discrimination in the somatosensory cortex. *Proc. Natl. Acad. Sci.*, 97:6091 – 6096.

Holmes, A. (1994). *Statistical Issues in functional imaging.* PhD thesis, University of Glasgow.

Jezzard, P., Matthews, P., and Smith, S. (2002). *Functional MRI: an introduction to methods.* Oxford Press.

Joachims, T. (2004). Multi-class support vector machine. http://svmlight.joachims.org/.

Kamitani, Y. and Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature neuroscience*, 8:679 – 685.

Kjems, U., Hansen, L., Anderson, J., Frutiger, S., Muley, S., Sidtis, J., Rottenberg, D., and Strother, S. (2002). The quantitative evaluation of functional neuroimaging experiments: Mutual information curves. *NeuroImage*, 15:772 –786.

LaConte, S., Strother, S., Cherkassky, V., and Hu, X. (2005). Support vector machines for temporal classification of block design fmri data. *NeuroImage*, 26:317 – 329.

Logotethis, N. and Wandell, B. (2004). Interpreting the bold signal. *Annu. Rev. Physiol.*, 66:735–769.

Mitchell, T., Hutchinson, R., Niculescu, R., Pereira, F., and Wang, X. (2004). Learning to decode cognitive states from brain images. *Machine learning*, 57:145 – 175.

Mountcastle, V., Steinmetz, M., and Romo, R. (1990). Frequency discrimination in the sense of flutter: Psychophysical measurements correlated with postcentral events in behaving monkeys. *The Journal of Neuroscience*, 10:3032 – 3044.

Mourão-Miranda, J., Bokde, A., Born, C., Hampel, H., and Stetter, M. (2005). Classifying brain states and determining the discriminating activation patterns: Support vector machine on functional mri data. *NeuroImage*, 28:980 – 995.

Nichols, T. and Holmes, A. (2001). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping*, 15:1 – 25.

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

Romo, R., Hernandez, A., Zainos, A., Brody, C., and Salinas, E. (2003). *Exploring the cortical evidence of a sensory-discrimination process*, chapter 5, pages 117–138. Oxford Press, first edition.

Salinas, E., Hernandez, A., Zainos, A., and Romo, R. (2000). Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *The Journal of Neuroscience*, 20:5503 – 5515.

Strother, S., Anderson, J., Hansen, L., Kjems, U., Kustra, R., Sidtis, J., Frutiger, S., Muley, S., LaConte, S., and Rottenberg, D. (2002). The quantitative evaluation of neuroimaging experiments: The npairs framework. *NeuroImage*, 15:747 – 771.

Vapnik, V. (1995). *The Nature of Statistical Learning Theory*. Springer Verlag.

Zigmond, M., Bloom, F., Landis, S., Roberts, J., Squire, L., and Woolley, R. (1999). *Fundamental Neuroscience*, chapter 26, pages 761–789. Academic Press, first edition.