

Anatomically Informed Basis Functions
—
Anatomisch Informierte Basisfunktionen

Dissertation

zur Erlangung des akademischen Grades

doctor rerum naturalium
(Dr.rer.nat.),

genehmigt durch
die Fakultät für Naturwissenschaften
der Otto-von-Guericke-Universität Magdeburg

von Diplom-Informatiker Stefan Kiebel
geb. am 10. Juni 1968 in Wadern/Saar

Gutachter: Prof. Dr. Lutz Jäncke
Prof. Dr. Hermann Hinrichs
Prof. Dr. Cornelius Weiller

Eingereicht am: 27. Februar 2001

Verteidigung am: 9. August 2001

Abstract

In this thesis, a method is presented that incorporates anatomical information into the statistical analysis of functional neuroimaging data. Available anatomical information is used to explicitly specify spatial components within a functional volume that are assumed to carry evidence of functional activation. After estimating the activity by fitting the same spatial model to each functional volume and projecting the estimates back into voxel-space, one can proceed with a conventional time-series analysis such as statistical parametric mapping (SPM). The anatomical information used in this work comprised the reconstructed grey matter surface, derived from high-resolution T1-weighted magnetic resonance images (MRI). The spatial components specified in the model were of low spatial frequency and confined to the grey matter surface. By explaining the observed activity in terms of these components, one efficiently captures spatially smooth response components induced by underlying neuronal activations localised close to or within the grey matter sheet. Effectively, the method implements a spatially variable anatomically informed deconvolution and consequently the method was named anatomically informed basis functions (AIBF). AIBF can be used for the analysis of any functional imaging modality. In this thesis it was applied to simulated and real functional MRI (fMRI) and positron emission tomography (PET) data. Amongst its various applications are high-resolution modelling of single-subject data (e.g. fMRI), spatial deconvolution (PET) and the analysis of multiple subject data using canonical anatomical bases.

Contents

Symbols and Abbreviations	6
1 Introduction	9
2 Basics	13
2.1 Magnetic Resonance Imaging	13
2.1.1 Spins in a magnetic field	13
2.1.2 Radio frequency pulses	14
2.1.3 Magnetic gradient fields	15
2.1.4 Measuring images	15
2.1.5 Parameters	17
2.1.6 Structural sequence	18
2.1.7 BOLD effect	18
2.1.8 Functional sequence	19
2.2 Voxel-based analysis	20
2.2.1 Preprocessing	21
2.2.2 Linear models	23
2.2.3 The univariate model	24
2.2.4 Statistical maps	25
2.2.5 Theory of Random Fields	26
3 Surface Reconstruction and Transformations	27
3.1 Voxels and Vertices	28

3.2	Surface reconstruction	28
3.3	Inflation and Flattening	30
3.4	Coregistration	30
3.5	Applications	31
4	Model specification	35
4.1	Introduction	35
4.2	The spatial domain	36
4.2.1	The spatial equation system	36
4.2.2	Specification of basis functions on the flattened cortex	37
4.2.3	Reprojection and transformation	38
4.2.4	Convolution by an additional point spread function	39
4.2.5	Estimating the spatial distribution of activity for one time point	39
4.2.6	Regularization	40
4.2.7	Global basis function	41
4.3	The temporal domain	42
4.3.1	Reprojection of parameters and deconvolution	42
4.4	Inferences about evoked responses over time	43
4.4.1	Univariate inferences	43
4.4.2	Multivariate inferences	45
4.5	Summary	46
5	Features of the method	51
5.1	Generation of simulated data	51
5.1.1	Functional magnetic resonance data	52
5.1.2	Positron emission tomography data	52
5.1.3	Ideal assumptions	53
5.2	Simulated studies	53
5.2.1	Sensitivity	54
5.2.2	Localization	57

5.2.3	Exploration of parameter space	61
5.2.4	Regularization	67
5.2.5	Spatial deconvolution	70
6	Applications	76
6.1	Functional magnetic resonance data	76
6.1.1	Single subject studies	77
6.1.2	Multiple subject study	82
6.2	Positron emission tomography	86
6.2.1	Multiple subject analysis	86
7	Discussion	90
7.1	Applications	90
7.1.1	High-resolution modelling	91
7.1.2	Group studies and sensitivity	92
7.2	Error sources and Limitations	94
8	Zusammenfassung	96
	Acknowledgements	99
	Curriculum Vitae	100
	References	102

Symbols and Abbreviations

MR Physics

M	magnetization
m	magnetization in rotating coordinate system
T	time
B	magnetic field strength
γ	Larmor constant
f	frequency
ω	angular frequency, $\omega = 2\pi f$
ϕ	phase
ρ	spin density
μ	magnetic permeability
χ	magnetic susceptibility
G	magnetic field gradient

General

N	quantity, number of
Y	observed data vector
B	spatiotemporal parameter matrix
β, γ	parameter vector
l	location
R	residual vector
I	identity matrix
P	residual forming projection matrix
ϵ	error vector
t	t-value
f	f-value
ν	degrees of freedom
σ	standard deviation
c, C	contrast vector, contrast matrix
Z	search volume
q, Q	CVA images, CVA image matrix

Temporal model

M	design matrix
K	convolution matrix

Spatial model

S	surface graph
V	vertex matrix
F	face matrix
f,g	arbitrary functions

L	spatial convolution matrix
\mathbf{b}	spatial basis function
M_G, G, A	spatial model matrices
w	width
λ	regularization factor
θ	hyper-parameters
W	weighting matrix
n	normalization function

Quantities

N_Y	number of images, scans
N_K	number of functional voxels
N_p	number of basis functions
N_{V_G}	number of grey matter vertices
N_W	number of regularization constraints

Subscripts

i	image, scan
j	parameter, basis function
l	voxel
G	grey matter surface
I	inflated surface
F	flattened surface
Y	in voxel-space
L	convolved in voxel-space
0	interface between white and grey matter

Abbreviations

MRI	Magnetic Resonance Imaging
fMRI	functional Magnetic Resonance Imaging
EPI	Echo Planar Imaging
TE	Echo Time
TR	Repetition Time
RF	Radio Frequency
BOLD	Blood Oxygen Level Dependent
HRF	Haemodynamic Response Function
PET	Positron Emission Tomography
SPECT	Single Photon Emission Computed Tomography
EEG	Electroencephalogram
MEG	Magnetoencephalogram
CBF	Cerebral Blood Flow
CSF	Cerebrospinal Fluid
FWHM	Full Width at Half Maximum

PSF
AR
CVA
TMS

Point Spread Function
Auto Regression
Canonical Variates Analysis
Transcranial Magnetic Stimulation

Chapter 1

Introduction

Over the last few years, *functional brain imaging* has become a rapidly developing part of neuroscience. Functional brain imaging involves the design of experiments followed by the acquisition and spatiotemporal characterization of functional brain data to characterize the underlying neuronal or metabolic events. Several modalities are used by researchers to acquire brain data, which are mainly functional magnetic resonance imaging (fMRI), positron emission tomography (PET), electroencephalogram/magnetoencephalogram (EEG/MEG), transcranial magnetic stimulation (TMS) and single photon emission computed tomography (SPECT). MRI and PET measure metabolic responses, whereas EEG/MEG and TMS measure neuronally mediated electrical and induced magnetic activity. This thesis is about the analysis of MRI and PET data in the context of functional brain imaging.

Many methods have been proposed for analyzing fMRI or PET data. Because experimentally controlled factors inducing signal changes over time are used to detect activations in voxel-space, the challenge is to find an appropriate temporal model to characterize the functional data. This temporal model can then be fitted at each voxel. These methods are called voxel-based methods and will be referred to as conventional methods in the remainder of this thesis. One common feature of many conventional methods is that they implicitly specify a spatial model by convolving the data with some spatial lowpass filter prior to the temporal analysis. Since the convolution kernel is the same at each voxel-position, the spatial model specified by this approach assumes spatial invariance with respect to the underlying tissue and cannot take into account the convoluted brain structure. One can apply conventional voxel-based methods without any prior filtering, but this also implies an invariant spatial model where all activation clusters have sub-voxel size. In both cases, the spatial models are spatially invariant and anatomically uninformed. After fitting the temporal model at each voxel, conventional methods have to solve the ensuing multiple comparison problem to make meaningful inferences about the data. This can be done by using results from Gaussian random field theory. Here, the resulting statistical maps are characterized over space, but because no anatomical information was introduced at an earlier processing stage, the inference is also anatomically uninformed. Voxel-based methods usually apply some masking of the functional data prior to the temporal modelling after smoothing to exclude voxels

at uninteresting positions, e.g. in the extracranial space. This binary masking should not be confused with a spatial modelling, because masking is just a device to exclude uninteresting time-series from the analysis.

There might be various reasons why it is not common practice to introduce spatial information into a functional analysis: (i) It is not yet really clear, how neuronal responses are related to the underlying neuroanatomy, (ii) an appropriate way to include anatomical knowledge into a spatiotemporal model may not be known, (iii) a spatiotemporal model would spoil the simplicity of temporal analyses, where one can analyze functional data independently of structural data, (iv) as it will be shown in this thesis, the specification and fitting of a spatial or spatiotemporal model is demanding in terms of computational resources and finally (v) a valid temporal model is an absolute necessity for proper characterization of the data, whereas one can analyze the data with a simple spatially invariant model.

In this thesis, a method is proposed to incorporate anatomical information into the analysis of functional data that resolves many of these problems. First it is assumed that such a model is separable in terms of its spatial and temporal components. The ensuing spatial model is based upon an approximation of the cortical grey matter sheet and uses a linearized approach to transform the functional data to a sparser representation. This new spatial representation enhances spatial components, which are assumed to carry interesting signal components and attenuates other components. These components are estimated for each time point and then analyzed over time, where one can use time-series analysis implemented in conventional methods. In the remainder of this thesis, the spatial model and the proposed method is called anatomically informed basis functions (AIBF).

The important point about AIBF is that it allows one to incorporate any anatomical knowledge, which can be described in the form of a set of spatial basis functions. In this work, the reconstructed cortical sheet was used as a basis for the spatial model. By making the assumption that the activations originate in the cortical grey matter sheet or from a location nearby, one can specify a model to capture spatial low-frequency components, which are considered to carry evidence of the blood oxygen level dependent (BOLD) effect (fMRI) or other metabolic responses (PET). There are many ways to model low frequency distributions in the two-dimensional cortical sheet embedded in the three-dimensional image space. Basis functions were used in this thesis, which can fit two-dimensional low-frequency components in each image. Two general classes of basis functions were employed. The first is a local basis function set, which captures a range of frequencies localized in space. The concept underlying this set of basis functions is similar to the Gabor transform as used in low-bandwidth compression of natural images. The other class of models we consider is essentially a global discrete cosine set encompassing the whole cortical sheet such that the distribution on the cortical sheet is decomposed into non-localized frequencies. This is similar to a Fourier transform.

The main aim of this work is to provide a proof of concept that AIBF can improve the characterization and interpretation of functional brain images. Three fundamentally different applications of AIBF are proposed. The first is concerned with high-resolution modelling of single subject fMRI data. It is foreseen that with increasing image reso-

lution in fMRI and accumulating knowledge about the relationship between MR measurements and underlying neurophysiology, spatially precise modelling will become a useful tool for the neuroimaging community. In this application, the focus lies on the anatomically informed, high-resolution of the resulting statistical maps.

The second application is the analysis of data, where the width of the image point spread function (PSF) is wider than the voxel-size. This is the case for PET and SPECT data, where AIBF can be used as a tool to implement anatomically informed spatial deconvolution.

The third application is the analysis of multi-subject data. The location of a functional area, which shows activation in response to a specific task, varies in a standardized anatomical frame over subjects. In conventional analyses, this variability is treated by convolution with a wide low-pass filter kernel such that the individuals' activations are smeared together. In other words, one trades resolution for sensitivity. With AIBF, one essentially does the same thing, but it is assumed that the anatomical variability consists of two components. One is the translational error tangential to the underlying cortical surface and the other is in the direction orthogonal to the cortical surface to provide a characterization that is much more precise in the orthogonal (normal) direction. By essentially projecting the data onto a smooth canonical grey matter surface the latter component can be removed from the data. It is shown in chapter 6 that this approach, combined with anatomically informed spatial deconvolution, offers advantages in comparison to the conventional way of analyzing group data.

The arrangement of the chapters is as follows: In the second chapter, some basics about MRI measurements are reviewed and the conventional voxel-based approach is outlined. The third chapter deals with the reconstruction of cortical surfaces. The fourth chapter is the core of the thesis. The spatial model is described and the mathematical foundation for specifying instances of the model is laid. The fifth chapter uses simulated activation data to explore the features of AIBF and it is shown that under ideal conditions AIBF provide better localization and increased sensitivity as compared to voxel-based methods. In the sixth chapter, AIBF is applied to real fMRI and PET data.

In more detail, chapter 2 deals with the basic background of this work and is divided into two sections. The first gives a brief overview of MR imaging, as it is typically employed in functional brain imaging experiments. The second describes a conventional voxel-based method (Statistical Parametric Mapping, SPM) to analyze functional data. Note that the temporal model used by SPM is also used to analyze the components extracted by AIBF.

The subject of chapter 3 is the surface reconstruction and flattening of an individual's brain. The described techniques have been implemented in various software packages. In this thesis, the Brainvoyager (Brain-Innovations, Rainer Goebel) software was used to extract surfaces and transform them to a flat map.

Chapter 4 is the core of this thesis. The framework of the spatial model that embodies anatomical information is introduced and described. Making an inference using AIBF consists of several processing stages. First, a spatial model based on the cortical grey

matter sheet is specified. This model is then fitted to each functional observation giving an estimated parameter matrix. After re-projecting these estimates back into voxel-space for each time point, one can analyze the projections over time and make inferences about the effects observed in image space (or indeed in the space of basis functions).

In chapter 5, AIBF is applied to simulated fMRI and PET data. This is done to validate the method, explore and discuss its salient features in relation to a conventional voxel-based method (SPM99). It is shown that AIBF offers, under ideal conditions, better resolution and/or sensitivity.

In chapter 6, AIBF is applied to real fMRI and PET data. Here, the conditions are less ideal, because various sources of error are introduced e.g. mis-registration of the structural and functional images, residual distortions in echo planar imaging (EPI), mis-estimation of the cortical surface, false assumptions about the distribution of activation on the cortical sheet and artefactual signal sources. Nevertheless, it is shown that AIBF is a useful tool to analyze functional data and can improve the resolution or sensitivity of the statistical tests.

In chapter 7, the method, its applications and potential pitfalls are discussed.

Finally, a summary in German is given in chapter 8.

Chapter 2

Basics

The basic background for the thesis is provided in this chapter. The first section deals with the concepts of magnetic resonance imaging (MRI) with respect to structural (T_1 -weighted) and functional (T_2^* -weighted) brain imaging. A fairly full account of the MR physics is provided, because both the functional data characterized and the anatomical information used by AIBF depend on MRI. However, the main focus of this chapter is the treatment of MRI data after acquisition.

After dealing with the physical basics of imaging data, the process of analyzing functional brain data with a univariate, voxel-based method is outlined. This is done based on the theoretical background and implementation of the *Statistical Parametric Mapping* software package in its current version, SPM99.

There are several reasons for this choice. SPM is essentially a package of documented, maintained and open-source routines, which are programmed in a modular and general way such that they could be adapted to other applications, as was done in this thesis. SPM has also become some kind of standard, in the neuroimaging community, for analysis of PET and fMRI data, i.e. it is well tested and offers a wide range of standard models for analysis of functional data. Additionally, its latest version, SPM99, provides an automated batch mode, which can be used to analyze functional data without user interaction. This feature was extensively used for analyzing numerous simulation studies.

2.1 Magnetic Resonance Imaging

2.1.1 Spins in a magnetic field

In the following, the behaviour of spins in the presence of a magnetic field is described in terms of a classical model. If protons are exposed to a homogeneous magnetic field with field strength B_0 [Tesla], they can have one of two possible states, a low or a high energy state. The directions of the spins are a function of their states and are parallel or antiparallel to the direction of the magnetic field. Additionally, in the

presence of the magnetic field, each spin precesses with the Larmor frequency, which is $f_L = \gamma B_0$, where γ is the Larmor constant and equals roughly 42.58 MHz/T. At magnetic equilibrium, at room temperature, there are more spins in the low than in the high energy state, which means that after reaching equilibrium there will be a longitudinal magnetization M_0 in the direction of the magnetic field. The process of getting to the magnetic equilibrium from a non-equilibrium state is characterized by two subprocesses: (i) *spin-lattice relaxation*: The spins exchange energy with the lattice (in the brain: liquid, water) until the equilibrium is reached. This process is governed by the longitudinal relaxation time T_1 . (ii) *spin-spin relaxation*: Spins experience, due to their surrounding spins, slightly different magnetic field strength such that their precessing frequencies vary over space on a microscopic level. This leads to a dephasing of spins (given that they were in phase initially, see next section) such that transverse magnetization decays over time. This kind of relaxation is parameterized by the transverse relaxation time T_2 . The dynamics of the magnetization components can be described by the Bloch equations:

$$\dot{M}_x = \gamma(\mathbf{M} \times \mathbf{B})_x - M_x/T_2 \quad (2.1)$$

$$\dot{M}_y = \gamma(\mathbf{M} \times \mathbf{B})_y - M_y/T_2 \quad (2.2)$$

$$\dot{M}_z = \gamma(\mathbf{M} \times \mathbf{B})_z - (M_z - M_0)/T_1 \quad (2.3)$$

where \times denotes the outer product of two vectors, \mathbf{M} is the magnetization vector and M_x, M_y and M_z are its components in x-, y- and z-direction.

2.1.2 Radio frequency pulses

The concept of acquiring MRI data is based on emitting radio frequency (RF) pulses in the presence of the magnetic field with strength B_0 in z-direction. The RF pulse will generate a magnetic field with strength B_1 , which rotates in the x-y-plane at frequency ω_0 such that the resulting magnetic field is given over time as

$$\mathbf{B}(t) = B_1 \cos(\omega_0 t) \mathbf{e}_x - B_1 \sin(\omega_0 t) \mathbf{e}_y + B_0 \mathbf{e}_z \quad (2.4)$$

where $\mathbf{e}_x, \mathbf{e}_y$ and \mathbf{e}_z are unit vectors in x-, y- and z-direction.

After transforming the Bloch equations to a rotating frame of reference, the magnetization in this rotating coordinate system is denoted by m_x, m_y and m_z . The Bloch equations in the rotating system are

$$\dot{m}_x = (\gamma B_0 - \omega_0) m_y - m_x/T_2 \quad (2.5)$$

$$\dot{m}_y = -(\gamma B_0 - \omega_0) m_x + \omega_1 m_z - m_y/T_2 \quad (2.6)$$

$$\dot{m}_z = -\omega_1 m_y - (m_z - M_0)/T_1 \quad (2.7)$$

where $\omega_1 = \gamma B_1$. Choosing $\omega_0 = \gamma B_0$ and assuming the actual duration of the RF pulses will be small compared to the relaxation times T_1 and T_2 , the equations simplify to

$$\dot{m}_x = 0 \quad (2.8)$$

$$\dot{m}_y = \omega_1 m_z \quad (2.9)$$

$$\dot{m}_z = -\omega_1 m_y \quad (2.10)$$

In summary, a short duration RF pulse as specified in Eq. 2.4, and emitted at the Larmor frequency of the precessing spins, generates a rotating magnetic field, which decreases the strength of the longitudinal z-component, but increases the strength of the transverse y-component as a function of the duration and strength of the RF pulse.

2.1.3 Magnetic gradient fields

To measure 2- or 3-dimensional images of some MRI parameter, one has to render the measured output a function of the location. This is done by applying magnetic gradient fields in addition to the homogeneous magnetic field B_0 . A gradient is defined by $G = \Delta B / \Delta x$, i.e. the field will change for a distance Δx by ΔB , where G is typically measured in $[mT/m]$. Note that the direction of the magnetic component always points in the same (z-)direction as the B_0 -component, but the magnetic field amplitude induced by G is a function of location. Since the Larmor frequency f_L is a function of the actual field strength, an applied gradient field changes the Larmor frequency as a function of location in the direction of the gradient. The Larmor frequency at location x is then given by

$$\omega(x) = \gamma x G \quad (2.11)$$

and the phase at location x and after time t is

$$\phi(x, t) = \gamma x \int_0^t G(\tau) d\tau \quad (2.12)$$

$$\text{or as a further approximation: } \phi(x, t) = \gamma G x t \quad (2.13)$$

2.1.4 Measuring images

In the following, the concepts of a 2-dimensional gradient-echo sequence will be described. First, three different uses of gradient fields are introduced.

A *slice-selective excitation* means a combination of a gradient and a RF pulse. First, a gradient field is applied in the z-direction such that the Larmor frequency of the spins is dependent on their position along the z-axis. Then a short RF pulse can be applied which effects only spins in a specific frequency range so that a transverse magnetic component is generated only for a confined x-y-slice. However, the gradient will rapidly dephase the spins because of the location-dependent Larmor frequencies. Therefore a negative gradient in the z-direction is applied after the first (positive) gradient to bring the spins back into phase.

A *phase-encoding gradient* is a gradient field applied in the y-direction. According to Eq. 2.12 the phase ϕ of the precessing spins is then a function of location in the y-direction and the gradient-time product.

A *frequency-encoding gradient* is a gradient field applied in the x-direction. The frequency ω of the spins will be a function of location (Eq. 2.11) in the x-direction.

After the slice-selective excitation a phase-encoding gradient G_P is applied. At the same time, the spins are dephased by a negative frequency-encoding gradient G_R , also called the *read-gradient*. The read-gradient is then inverted and after some time the spins are back in phase and a gradient echo signal is formed, which is measured over time. This process is iterated for different values of G_P ranging from $-G_{P_0}$ to G_{P_0} . For a single iteration, let t_P be the duration of G_P , then the phase of the spins at a position (x, y) during acquisition of the signal is given by

$$\phi = \gamma G_P y t_P \quad (2.14)$$

and their frequency by

$$\omega = \gamma G_R x \quad (2.15)$$

meaning that the signal measured at (x, y) is

$$\exp(i(\omega t + \phi)) = \exp(i\gamma G_R x t) \exp(i\gamma G_P y t_P) \quad (2.16)$$

Let $\rho(x, y)$ be the spin density in the excited slice, then the total measured signal for the slice for one phase-encoding step is

$$S(t, G_P) = \int_x \int_y \rho(x, y) \exp(i\gamma G_R x t) \exp(i\gamma G_P y t_P) dx dy \quad (2.17)$$

After defining $k_x = \gamma G_R t$ and $k_y = \gamma G_P t_P$, this can be written as

$$S(k_x, k_y) = \int_x \int_y \rho(x, y) \exp(ixk_x) \exp(iyk_y) dx dy \quad (2.18)$$

which is the Fourier transform of the spin density. By retransformation one can compute $\rho(x, y)$.

The space, in which the signal is measured at discrete coordinates (k_x, k_y) is called the K-space. The resolution in retransformed image space depends on the gradient-time product. This can be seen by transforming the definition of a gradient:

$$\Delta x = \frac{2\pi}{\gamma G_R t_A} \quad (2.19)$$

$$\Delta y = \frac{2\pi}{\gamma G_P t_P} \quad (2.20)$$

where t_A is the time from the centre of the excitation pulse to the end of acquisition. In other words, there are theoretically two ways to increase the resolution in image space. The first is to prolong times t_P and t_A and the second is to apply stronger gradients.

2.1.5 Parameters

Longitudinal relaxation time T_1

The following observation is the basis of measuring the distribution of T_1 in a sample. Without the influence of magnetic fields other than B_0 the magnetic component M_z approximates the magnetic equilibrium M_0 in the z-direction over time as characterized by

$$\dot{M}_z = (M_0 - M_z)/T_1 \quad (2.21)$$

as can be seen from Eq. 2.3. If one defines the initial state $M_I = M_z(0)$, the solution of the equation yields

$$M_z(t) = M_0 + (M_I - M_0) \exp(-t/T_1) \quad (2.22)$$

Assuming that one can invert the equilibrium magnetization M_0 in the z-direction at $t = 0$ such that $M_z(0) = -M_0$, then

$$M_z(t) = M_0(1 - 2 \exp(-t/T_1)) \quad (2.23)$$

The inversion of M_0 can be done by application of a 180° pulse. A suitable sequence for measuring T_1 is to apply a 180° pulse, wait time T_w , apply a 90° pulse and measure the resulting signal. To quantify T_1 , one has to sample Eq. 2.23 several times for a range of T_w and find a monoexponential fit with parameters T_1 and M_0 . Note that, if $T_w = T_1 \ln 2$, $M_z(T_w)$ vanishes and T_1 is known. However, this process of quantifying T_1 takes a rather long time due to the need to measure several T_I and is not necessary for measuring a brain image which reflects the differences in T_1 of the tissues. In general, such a T1-weighted image is measured by sampling Eq. 2.23 once. Further details about the T1-weighted sequence used in this thesis is provided in 2.1.6.

Effective transverse relaxation time T_2^*

The transverse relaxation time T_2 is assumed to be related to two underlying effects. The first is the spin-lattice relaxation, which is an energy transfer of the spins and results not only in an increase of longitudinal, but also a loss of transverse magnetization. The other effect is related to the microscopic inhomogeneities of the magnetic field, resulting in a dephasing of spins due to different precessing frequencies. It follows directly from this that T_2 must always be shorter than T_1 , since the dephasing of spins affects only T_2 .

Additionally, due to local field inhomogeneities on a macroscopic level, spin assemblies experience different magnetic fields and dephase even faster. The transverse relaxation time associated with this effect is T_2' . The effective transverse relaxation time is given by

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \quad (2.24)$$

The magnetic field inhomogeneities associated with T_2' are assumed to be constant over (experiment) time, whereas the microscopic spin-spin relaxation effects associated with T_2 are not. This means that signal loss due to T_2' can be recovered by a spin-echo sequence and T_2 without the T_2' contribution can be measured (described in e.g. (Haacke et al., 1999)).

A T_2^* -weighted image can be measured in principle by a gradient-echo sequence as described in section 2.1.4.

2.1.6 Structural sequence

In this thesis, two different sequences were used to acquire structural data. Most of the T1-weighted anatomical images were acquired on a clinical 1.5 T Siemens Vision unit (Siemens GmbH, Erlangen) with a 3D FLASH-sequence with repetition time (TR) = 15 ms, echo time TE = 5 ms, flip angle = 30° . Each image matrix consisted of 256×256 voxels, 192 partitions, sagittal orientation, voxel size $1 \times 1 \times 1 \text{ mm}^3$. The repetition time TR is the time between two RF pulses. The echo time TE in a gradient-echo sequence is the time between RF-pulse and the time point when applied negative and positive gradient lobes add to zero gradient-time product.

Another set of T1-weighted images was acquired using the inversion recovery sequence as described in (Deichmann and Turner, 2000).

2.1.7 BOLD effect

The BOLD (blood oxygen level dependent) effect is used to detect differences or changes in local oxygenation states of the blood in the brain.

It is known that deoxygenated blood is slightly more susceptible than oxygenated blood, where the susceptibility is defined as

$$\chi = \frac{\mu}{\mu_0} - 1 \quad (2.25)$$

where μ_0 is the permeability of empty space and μ the permeability of a location in an object. For diamagnetic material, $\mu < 1$ and for paramagnetic material $\mu > 1$ such that χ is negative for diamagnetic and positive for paramagnetic material.

A susceptibility gradient along brain tissue causes a variation in the local magnetic field, which in turn leads to different Larmor frequencies along this susceptibility gradient. This is measured as a signal loss in a T_2^* -weighted sequence due to a shorter T_2' .

In the brain, oxygenated blood in vessels have a similar susceptibility to the surrounding tissue, whereas deoxygenated blood is more paramagnetic so that there is a measurable signal loss around blood vessels due to a susceptibility gradient.

Neural activity triggers a vascular mechanism that supplies the area of neural activity with fresh oxygenated blood to compensate for the increased metabolic demand for oxygen. However, this supply is higher than the actual uptake. Although there is an

increase of deoxygenated blood due to the uptake, the increased flow renders the blood more oxygenated as compared to the state prior to the neural activity. This process results in a signal increase during the increased flow. Parameterized models of the haemodynamic response function were proposed by (Friston et al., 1994b; Lange and Zeger, 1997; Cohen, 1997). In Fig. 2.1, a model of the BOLD effect consisting of two parameterized Gamma functions is shown.

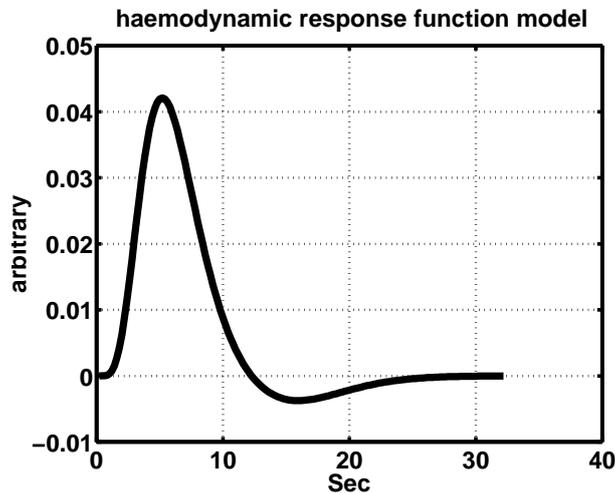


Figure 2.1: Model of haemodynamic response function consisting of the sum of two gamma functions

Note that the typical BOLD response curve not only shows a latency, but also a dispersion of some seconds. It was shown by Aguirre et al. (1998) that the exact shape of the BOLD effect in a given brain area varies across subjects, but there are less variations intra-subject over time. In (Kastrup et al., 1999), it was observed that the hypercapnia-induced BOLD response varies across brain regions, reflecting regionally specific influence of factors such as cerebral blood flow, volume, oxidative metabolism and vascular anatomy.

2.1.8 Functional sequence

Echo planar imaging (EPI) is, because of the speed of acquisition, the most often used sequence class to measure $T2^*$ -weighted images, i.e. evidence of the BOLD effect. The important point about an EPI sequence is that one RF pulse is used to sample multiple echos. In a typical two-dimensional gradient-echo EPI sequence, the aim is to cover all K-space positions in a slice by using one RF pulse followed by rapid x- and y-gradient switching. The z-gradient would select the slice, followed by a strong negative phase-encoding gradient and switching on of a read-gradient (frequency-encoding) during sampling one line in K-space. Between lines, a positive blipped phase-encoding gradient is applied to prepare sampling the next line in K-space while switching on the inverted read-gradient. This procedure is repeated with alternating read-gradients until all positions in K-space for the given slice are covered. Clearly, because $T2^*$ is around 60

ms, TE should be around this time. This imposes some constraints on the gradient-hardware used. For instance, one constraint is that the read-gradient must be able to alternate, in a short time, between positive and negative values.

The functional data were acquired on a clinical 1.5 T Siemens Vision unit (Siemens GmbH, Erlangen) with TR = 168 ms, TE = 79 ms, flip angle = 90^0 . Each image consisted of 128×128 voxels, 20 slices, transverse orientation, voxel size $1.8 \times 1.8 \times 3$ mm³. Note that in a two-dimensional EPI sequence with enough slices covered, the flip-angle can be 90^0 , because the time between RF pulses of the same slice is long.

A disadvantage of EPI, which is of particular importance in this work, is that EPI data is subject to geometrical distortions and, like other $T2^*$ -measuring sequences, to signal dropouts. This is caused by local field inhomogeneities existing at tissue interfaces. Typically, in functional images such local inhomogeneities and geometric distortions can be observed around the petrous bone and the frontal sinuses. Local field inhomogeneities cause local magnetic gradients resulting in two types of artefacts. The first is a signal dropout, i.e. a decrease in signal intensity, in the presence of local magnetic gradients, which cause a complete dephasing of spins such that no gradient echo can be measured in these regions. This artefact is caused by a long TE, i.e. the more time elapses between RF pulse and gradient-time product cancelling, the more dropout occurs. Dropout effects can be reduced by decreasing TE, but the BOLD effect cannot be detected well at a shorter TE than say 30 ms. The other type of artefact is geometric distortion along the phase encoding direction when using a trapezoidal K-space sampling as described above. This is caused by the fact that neighbouring points in K-space in the x-direction are measured with minimal delay such that nearly no additional dephasing of spins can occur. However, neighbouring points in K-space in the y-direction are measured with a much longer time interval such that more dephasing can occur and distortions in the phase-encoding direction can be observed.

Post-processing schemes based on field maps can be employed to estimate an image transformation to correct geometric distortions ([Jezzard and Balaban, 1995](#)).

2.2 Voxel-based analysis

In this section, the concepts of univariate, voxel-based analysis of functional data in voxel-space are described. Generally, any analysis process can be subdivided into three parts: preprocessing, model estimation and statistical inference.

In the preprocessing stage, the raw data (reconstructed MR or PET images) are transformed such that the coordinate systems of all images are aligned to the same anatomical space. This has the effect that a set of transformed image coordinates describes the same anatomical location in each image. The preprocessing stage also includes any spatial filtering.

In the modelling step, all prior temporal knowledge or assumptions about the expected response are summarized within a linear model (described by a user-specified design matrix) and the data is fitted, in a least-squares fashion, to the model yielding an

estimated parameter vector at each voxel.

The inference step assesses using an F-test, whether a projection of the data into a subspace, spanned by a particular set of linear functions of the design matrix columns, is large compared to some measure of the residual variance. The F-ratio is used to reject the null hypothesis that the additional compound of regressors do not explain more of the data than the reduced model. As a special case of an F-test, one can also employ a t-test, where the projection subspace is one-dimensional. The F or t statistic for each voxel constitutes a statistical parametric map, or SPM, upon which inference is based.

2.2.1 Preprocessing

In this subsection, the concepts of the necessary preprocessing steps prior to the analysis of functional data are described. The overall goal of the various preprocessing steps is to ensure that the voxel-space of each image relates to the same anatomical space. This is accomplished by registering the images together using affine and non-linear transformations. In neuroimaging, one can differentiate between two cases of registering a pair of images: (i) Intra- or inter-modality intra-subject registration and (ii) inter-subject registration.

Intra-modality intra-subject registration

Subject movements have to be expected, even if the subject is highly motivated and instructed to lie still. Subject movements tend to be between 0.5 and 1 mm, where the amplitude of movements usually increases with the duration of experiment. The main reason for applying some kind of post-hoc registration is to remove motion artefacts from the data and thereby increase the sensitivity of the analysis. For example, it is quite common to observe a linear translation of the subject's head in the z-direction over time. If there is an underlying activation, the subject's movement would apparently smear this activation in the z-direction. By registering all images to a common spatial frame, this smearing effect is largely removed and, depending on the shape of the activation, the sensitivity of the analysis with respect to the activation is increased. The most commonly used approach is to assume that a six parameter rigid-body transformation of each image to a common spatial frame solves the realignment problem. Such a transformation can be found e.g. by minimizing the sum of squares of the image difference between a pair of images (Woods et al., 1992; Friston et al., 1995a). One way to do this is via a Gauss-Newton optimization as implemented in SPM. Note that there are several reasons why a rigid-body transformation can only be an approximation for fMRI data. Two important issues are that (i) spatial distortions in fMRI images are partially a function of location within the magnetic field such that large subject movements can locally change brain shape (Jezzard and Clare, 1999) and (ii) a fMRI volume is acquired as a set of slices, so movement during acquisition of a volume introduces non-rigid transformations. In PET, large movements can render the attenuation correction incorrect (Andersson et al., 1995).

Inter-subject registration

Registration of different subjects' brain images is an important concept in neuroimaging. It is used to transform images from different subjects to the same anatomical standard space, e.g. the space defined in Talairach and Tournoux (1988). Its application is critical in the analysis of functional multiple-subject data and reporting results within a standardized frame. An incidental but interesting analysis is to use information derived from the transformations to make inferences about differences in brain shape over subjects (Ashburner et al., 1998; Gaser et al., 1999; Ashburner and Friston, 2000). With respect to functional studies, the definition of a goal function of the inter-subject registration is difficult, because it is known that a one-to-one registration of functional anatomy does not exist between any two brains. Therefore, a typical assumption is that some residual variability cannot be removed from the data even after applying an inter-subject registration. This residual anatomical uncertainty is then resolved by applying a spatial smoothing operation prior to the estimation procedure (see section 2.2.1). This is why constraints for inter-subject registration can be of a low-frequency nature (with respect to the resolution of the functional data), i.e. the transformation used for functional analysis does not need to incorporate high-resolution constraints. In SPM, inter-subject registration is parameterized by a twelve-parameter affine transformation, followed by a nonlinear transformation with a low-frequency discrete cosine set (Ashburner and Friston, 1999). Other commonly used methods based on multiscale or polynomial basis functions are (Collins et al., 1994; Woods et al., 1998).

Inter-modality intra-subject registration

For different sequences or modalities (e.g. functional MRI, structural MRI, PET), the image intensity distributions for the brain tissues are different. Therefore intensity difference minimization as applied in inter-modality registration is not an appropriate goal function to find a transformation. Instead, one can reduce the problem to a mixture of inter-subject registration and segmentation problems (Ashburner and Friston, 1997). First, one uses the inter-subject registration to transform the target to a template image of the same modality, where this registration is a twelve-parameter affine transformation subject to a minimization of a function of image intensity differences. Given a probabilistic map of the brain tissue distribution in the template images, one can segment the original images based on this information. Once the three major brain partitions, grey matter, white matter and CSF are estimated, these can be registered separately using both template and target segments. The final transformation is then given by the mean of the transformations of the three partitions. Note that functional EPI images can be locally distorted with respect to the underlying neuroanatomy. This problem is of particular importance for this thesis and will be treated in more detail in section 3.4.

Spatial smoothing

Having transformed all images to the same spatial frame, the decision must be made about whether to apply a spatial smoothing filter to the data prior to the statistical analysis. Generally, a smoothing filter will decrease the spatial resolution of the data while increasing the sensitivity to signal sources of the same width as the filter kernel (Friston et al., 1996; Kruggel et al., 1999).

Given that one generally employs a conventional analysis in voxel-space, it is a good idea to apply a smoothing filter, because the measured BOLD-effect can have an intrinsic width of several millimetres, whereas the voxel-size in fMRI is usually smaller than this. A smoothing filter of the same width as the underlying signal (comprising several voxels) will enhance the signal-to-noise ratio. Activations can then be detected over time by a voxel-based approach, which will be described in the next section.

Accepting the usefulness of spatial smoothing, the question is, which filter should be chosen. Since the underlying, presumably non-stationary, spatial haemodynamic response is not known, one generally employs a spatially invariant, three-dimensional lowpass-kernel (e.g. Gaussian) with a full width at half maximum (FWHM) of two to three times the voxel-size in each direction (i.e. 4 to 8 mm).

Another reason to smooth the data with a spatial low-pass filter kernel is the need to make adjustments to the statistical inferences about the data. As described in Section 2.2.4, the theory of random fields can be employed to assess the significance of observations, but one of the necessary assumptions is that the analyzed spatial volume is a good lattice representation of a smooth continuous random field, which is approximately true after the smoothing operation.

As mentioned in section 2.2.1, smoothing sacrifices resolution to reduce the negative side effects of residual anatomical and functional variability so that smoothing appears to be mandatory for multiple-subject studies.

The important question, which spatially invariant smoothing kernel should be used, cannot be answered without having assumptions about the size of the underlying activations. Since these are presumably unknown and activations of different widths might co-exist, the usual approach is to choose the smoothing filter by experience, i.e. choose a filter that produced sensitive results in a former study. Another technique is to use a subset of the data to explore the space of smoothing kernel parameters and choose the filter with the best performance. In this context, multi-scale space techniques (Worsley et al., 1995) have been proposed to detect signals with smoothing kernels of different FWHM in the framework of four-dimensional random fields.

2.2.2 Linear models

It is apparent from the literature that linear models dominate functional brain imaging. At first glance, this is unexpected, because the human brain is a complex system exhibiting extremely non-linear features, presumably at any level of description. The important point to note is that methods like PET and fMRI merely show the noisy

image of blurred metabolic responses due to underlying neuronal activations at a rather macroscopic level. Therefore it is not really surprising that linear models have proved to be appropriate for most studies. They are robust and are good at describing smooth spatiotemporal signal sources measured at a low signal-to-noise ratio. Also, by using non-linear basis functions, the linear model can be extended to model non-linearities.

Basically, linear methods can be described as a projection to a subspace of the data space. This subspace is spanned by the columns of the so-called *design matrix*, where a column is called a regressor, explanatory variable or a *basis function*. Basis functions as used here do not need to be orthogonal or orthonormal, although this enhances the interpretability of the fitted effects.

In functional brain imaging, the usual way to analyze a set of data (consisting of a time-series of three-dimensional images) is to fit the time-series at each voxel with a linear combination of the basis functions, resulting in a vector of estimated parameters. One then makes inferences about these estimates using hypothesis tests. The design matrix is the same at each voxel location.

The specification of the design matrix is the critical modelling task. A neuroimaging experiment usually consists of a sequence of well-defined experimental conditions, where the basis functions are chosen such that the expected experimentally induced effects over time can be captured for each condition.

In the following, the key equations of the general linear model and subsequent inference are outlined.

2.2.3 The univariate model

In the remainder of this section, the subscript \mathbf{l} denoting voxel location is omitted, because the model and related equations are the same at each voxel. The general linear model at each voxel position is given by

$$\mathbf{Y} = M\beta + K_I\epsilon \quad (2.26)$$

where Y is the observed image intensity vector for images $i = 1, \dots, n$, M is the $n \times p$ -design matrix of the model, β a p -dimensional parameter vector, ϵ is a vector of errors, normally distributed with $\epsilon \sim N(0, \sigma^2 I)$ and K_I is the intrinsic temporal correlation matrix such that the serially correlated errors $K_I\epsilon \sim N(0, \sigma^2 K_I K_I^T)$.

For PET, K_I is assumed to be I_n , the identity matrix. For fMRI, the measurements are serially correlated. In the typically employed regime of two to five seconds for whole brain measurement, K_I must be estimated to assess the variance of the weighted estimated parameters $\hat{\beta}$ (see below). Estimating K_I allows pre-whitening of the data by pre-multiplication of K_I^{-1} and subsequent estimation of β . This is the most efficient (Worsley and Friston, 1995; Friston et al., 2000) estimator, but can lead to biased estimates of σ^2 , if the assumptions about K_I are wrong.

Another solution, adapted by SPM, is to multiply Eq. 2.26 by K_E , where K_E is some

convolution matrix. This changes the model to

$$K_E \mathbf{Y} = K_E M \beta + K_E K_I \epsilon \quad (2.27)$$

The autocorrelation matrix is then given by $K = K_E K_I K_I' K_E'$. A wrong assumption about K_I still biases the estimate of σ^2 , but depending on the form of K_E , the bias is attenuated at the cost of efficiency (Friston et al., 2000). In SPM, the most often chosen K_E is a bandpass filter such that low-frequency drifts and high frequencies above the expected spectrum of the haemodynamic response function (HRF) are removed from the time-series. K_I can be estimated using an $AR(1)$ -model or simply assumed to be I_n , given that $K_E K_I K_I' K_E' \approx K_E K_E'$.

Let $(K_E M)^+ = ((K_E M)' K_E M)^{-1} (K_E M)'$ be the pseudoinverse of $K_E M$, then the residual forming projection matrix P is given by (Rao and Toutenburg, 1995)

$$P = I_n - K_E M (K_E M)^+ \quad (2.28)$$

After estimating β by the least-squares method, the residuals are given by

$$\mathbf{R} = K_E (\mathbf{Y} - M \hat{\beta}) = P K_E Y \quad (2.29)$$

and R_i are the residual fields consisting of $R_i(\mathbf{l})$ for all voxel positions \mathbf{l} .

The variance σ^2 is estimated by dividing the sum-of-squares of the convolved residuals by its expectation, i.e.

$$\hat{\sigma}^2 = \frac{\mathbf{R}' \mathbf{R}}{\text{trace}(PK)} \quad (2.30)$$

2.2.4 Statistical maps

The estimated parameters are simply

$$\hat{\beta} = (K_E M)^+ K_E Y \quad (2.31)$$

Inferences about contrasts $\mathbf{c}' \hat{\beta}$ are made using the t-statistic

$$t = \frac{\mathbf{c}' \hat{\beta}}{(\mathbf{c}' \hat{\sigma}^2 (K_E M)^+ K (K_E M)^+ T \mathbf{c})^{1/2}} \sim t_\nu \quad (2.32)$$

where t_ν denotes a t-distribution with ν degrees of freedom. The hypothesis tested is that $\mathbf{c}' \beta = 0$, where \mathbf{c} is a suitable contrast vector (Frackowiak et al., 1997) and

$$\nu = \frac{\text{trace}(PK)^2}{\text{trace}(PKPK)} \quad (2.33)$$

are the effective degrees of freedom (Worsley and Friston, 1995).

2.2.5 Theory of Random Fields

The null-hypothesis is that $\mathbf{c}'\beta = 0$. Any effect that is captured by fitting the basis functions can cause $\mathbf{c}'\beta$ to deviate from zero and would cause a negative or positive shift of the t-value. By computing the probability $p(t)$ that a given t-value $t(\mathbf{l})$ is generated by a t-distribution with ν degrees of freedom, one can reject the null hypothesis at a voxel, if $p(t(\mathbf{l})) \leq \alpha$, where α is a specified error level. This scheme is a classical t-test and operates on the improbability (the probability of random occurrence) of a single observation. The same principle applies to an F-map, where one tests a whole subspace spanned by the basis functions weighted by a contrast matrix C .

In a spatially correlated statistical map, the null hypothesis $\mathbf{c}'\beta_{\mathbf{l}} = 0$ is tested at each location $\mathbf{l} \in Z$, where Z is the search volume. It is obvious that here the probability of observing a given t-value at any voxel location \mathbf{l} should also be a function of the number of observations and the level of autocorrelation. The so-called multiple comparison problem is to assess the probability of a given observation in a statistical map encompassing a large number of other observations.

To solve this problem, one could assume independence of t-values in the statistical map and apply a Bonferroni correction, but because of the spatial correlation, this method would yield very conservative probabilities of random occurrence (Holmes, 1994).

The method usually chosen to solve the multiple comparison problem is to treat the statistical map under the null hypothesis as being generated by random fields so that one can estimate the probabilities of observed events.

Worsley et al. (1996) show how the Euler characteristic of various statistical fields can be used to approximate, for high thresholds, the probability that a given maximum in a search volume exceeds a threshold t_{thres} . The search volume of the statistical map can have any shape, although the best approximation will be achieved, if the search volume is convex. The theory provided in Worsley et al. (1996) can also be used to estimate p-values for peak height in search volumes with a dimension smaller than three. This is important, if one wants to infer about statistical maps on a two-dimensional cortical surface. The approximation of the p-value for peak height in a search volume can be combined with results derived in Friston et al. (1994a) to estimate the probability of the size of a cluster that exceeds a user-specified spatial extent threshold.

The application of Gaussian field theory does not change the SPM but simply adjusts p-values to protect against family-wise false positives over Z , the search volume. The adjustment is a function of the spatial smoothness of the SPM component fields (usually approximated by the residual fields in Eq. 2.29).

Chapter 3

Surface Reconstruction and Transformations

In this chapter, the reconstruction and various transformations of a subject's grey matter surface are described. Additionally, some applications based on the reconstructed surface are discussed.

The aim of a reconstruction is to find a surface representation in three-dimensional space, which best approximates the unknown grey matter surface. The reconstruction is typically based on a subject's high-resolution T_1 -weighted MR image.

From a historical point of view, computerized surface reconstruction was motivated by manual cortical flattening approaches used for the study of visual areas, e.g. (van Essen and Zeki, 1978). With developing computing power, computerized surface reconstruction based on high-resolution T_1 -weighted MR imaging became feasible. Algorithms for flattening polyhedral surfaces were proposed by Schwartz and Merker (1986) and Schwartz et al. (1989). More holistic approaches solving the problem of transforming a given T_1 -weighted MR image to a reconstructed and flattened surface were proposed by Dale and Sereno (1993), Carman et al. (1995), van Essen and Drury (1997), Dale et al. (1999) and Fischl et al. (1999b).

As described in these papers, after reconstruction, the surface can be transformed to other less folded representations by decreasing local curvature at the costs of cuts and/or distortions. In particular, the focus in this thesis lies on the flattened cortical surface, i.e. the folded grey matter sheet projected to a plane. Other possible representations are the inflated surface and projections onto spheres or ellipsoids. In functional brain imaging, these projections are mostly used for visualization of functional data or for detection of shape differences over subjects. In this thesis the projections define a manifold that contains the AIBF that are used to characterize the functional data per se.

3.1 Voxels and Vertices

The representation of images in three-dimensional voxel space is a commonly used concept in neuroimaging, because it is a natural way to store images aligned on a regularly sampled grid. A voxel is sometimes described or understood in neuroimaging as something that represents a small cuboid volume. This view is convenient, but slightly misleading, because one effectively only deals with a discretely sampled measurement of a volume, i.e. a grid of measurements. The latter view is largely compatible with the first description, but has the advantage that a measured image volume can be seen as a graph with vertices (nodes), edges and faces, where the volume itself does not play a role. In such a voxel-space the positions of the vertices are ordered on a regular grid such that all vertex coordinates, the set of edges and faces are implicit in the specification of the three dimensions of this graph, i.e. one does not need to store all this information explicitly. One only needs to record the intensities measured at the vertices. The disadvantage of this representation is that the sampling of the grid is rigid, i.e. one cannot locally increase the resolution of the graph or remove irrelevant information by pruning vertices. The alternative to a graph specified in voxel-space is a graph where one explicitly specifies vertices, faces and/or edges. Here one has the advantage that one only describes relevant information, but at the price of storing additional information about vertices and faces. In the following, such a representation will be called a vertex-space representation, whereas the rigid-grid representation will be called voxel-space.

The cortical surface can be approximated both in voxel- and in vertex-space. A description of this surface in voxel-space is e.g. given by the segmentation of the grey matter sheet. However, this representation is only a coarse approximation to the smooth underlying extensively folded surface. The cost in space of increasing the resolution in voxel space is prohibitive and therefore surface reconstruction techniques use a first segmentation in voxel space to generate a graph in vertex space. This new representation can then be modified by employing a smoothness constraint to minimize the discretization error, evident in the jagged segmentation in native voxel-space.

In vertex space, the surface is described by a graph $S = (V, F)$, where V is a $N_V \times 3$ -matrix and F is a $N_F \times x$ -matrix, N_V and N_F are the number of vertices and faces, x is 3 or 4, i.e. faces are triangles or quadrangles. Such graphs can provide good approximations to smooth surfaces with minimal discretization error. In practice, the visualization of triangular meshes can be efficiently implemented by using an Application Programming Interface (API) like OpenGL.

Throughout this work, we used the software package BrainVoyager (version 3.8, Brain Innovations, Rainer Goebel), which deals with triangle-based representation of surfaces.

3.2 Surface reconstruction

For surface reconstruction, a structural MRI-sequence as described in Section 2.1.6 is used.

A surface reconstruction has two aims: (i) to identify a surface that best approximates the cortical grey matter and (ii) topology, i.e. connectivity along the surface.

The basic strategy, to reconstruct an individual’s surface of the cerebrum, consists of three steps: (i) segment the white matter volume of each hemisphere, (ii) triangulate the outer (closed) white matter surface of each hemisphere and (iii) adjust the surface by using an iterative intensity-based optimization routine.

For simplicity, later stages of the reconstruction process require one filled white matter partition of each hemisphere. This is done by removing the cerebellum and declaring ventricles and internal grey matter structures (e.g. thalami) as white matter. The surface of this white matter partition is triangularized by identification of the outside voxel faces and partitioning each of these faces into two triangles. This gives the surface $S_0 = (V_0, F)$, where V_0 is an $N_{V_0} \times 3$ -matrix of vertex coordinates, F is an $N_F \times 3$ -matrix of vertex indices, N_{V_0} and N_F are the number of vertices and faces. An important feature of this approach is that only those vertices are connected (part of the same face), which are also adjacent in terms of the underlying true cortical surface, i.e. the topological structure of the cortical surface is correctly mapped to S_0 . This holds for most brain volumes, because the minimal distance between the grey-white interface of the banks of a sulcus is typically larger than the structural resolution of the T_1 -weighted MRI. However, there are cases when this approach fails, especially in brain regions with a rather thin grey matter sheet like the occipital lobe. Here, *topological defects* of the resulting surface S_0 can sometimes be observed, e.g. vertices, which are close to each other in Euclidean image space, but remote from each other on the brain surface, are part of a face. Currently, no sophisticated algorithm is known to automatically detect and remove topological defects. This means the user must manually intervene, i.e. one identifies the source of error, mostly errors due to partial volume effects, and changes the white matter segmentation by removing (or adding) voxels.

Since the underlying grey matter surface is smooth (i.e. the local curvature is bound by some finite value), the first rough approximation S_0 is improved by smoothing each vertex coordinate with respect to its local vertex neighbourhood such that V_0 is transformed to a smooth representation of the white matter surface $S_W = (V_W, F)$. In a next step, a surface lying within the grey matter sheet is identified by translating the vertices in V_W using a local iterative intensity-based optimization routine, which gives a representation of the pial surface $S_G = (V_G, F)$.

An important technical point to note here is that all brain surfaces derived from a structural T_1 -weighted MR image are a modified version of S_0 , in the sense that the face coordinates matrix F never changes, only the vertex coordinates matrices V_W and V_G are modified to describe different surfaces. In this way, the coordinate transformation between the different surfaces is implemented by an exchange of vertex coordinates matrices.

It would be desirable to estimate a surface, which approximates a given grey matter layer, e.g. layer IV. However, this is not feasible due to the structural MRI resolution [1 1 1] mm, which is high enough to find a surface approximation within the grey matter sheet, but not within a certain layer. However, it is possible to find rather accurate

estimates of the white-grey interface and the pial surface due to the intensity differences between tissues and CSF. Another interesting application is then to estimate the local thickness of the grey matter sheet (Fischl and Dale, 2000; Jones et al., 2000; MacDonald et al., 2000).

3.3 Inflation and Flattening

To compute an inflated version $S_I = (V_I, F)$ of the grey matter sheet S_G , the vertex coordinates V_G are locally smoothed such that local curvature is reduced. To prevent serious local distortions, the optimization is subject to the further constraint that the original geometrical distances between adjacent vertices in S_G should remain the same. By user-dependent regularization between those two constraints, the result is an inflated version of S_G , in which both sulci and gyri are fully visible.

After placing several cuts on the medial side of a hemisphere, the inflated surface is projected onto the y-z-plane. Surface parts on the former medial side are folded outside such that the whole brain surface (except for the cuts) is visible in the y-z-plane. This first step guarantees that the starting condition of the following optimization necessary to remove residual distortions will not end in a minimum associated with a mirror-image configuration (Fischl et al., 1999b). Since a simple projection of a highly folded structure to a plane implies serious distortions, these are largely removed by a subsequent iterative optimization with the two constraints (i) the surface is kept in the y-z-plane and (ii) the lengths between adjacent vertices as measured in S_G are restored. However, there is no isometric mapping, i.e. one cannot obtain a flat and undistorted cortical map, because the Gaussian curvature of the convoluted brain differs from a surface in a plane. The resulting flat map $S_F = (V_F, F)$ is a locally distorted representation of the grey matter sheet S_G . The distribution of distortions depends on the location of cuts made to enable the surface unfolding. In the example shown below, five cuts were used.

The transformation from S_G to S_I is implemented by an exchange of the vertex coordinate matrices V_G with V_I . In the case that a surface contains only a subset of the original vertices and faces, this information must be stored either in the form of a binary visibility or an index reference vector with respect to the full surface.

3.4 Coregistration

As pointed out in section 2.2.1, an accurate co-registration between functional and structural data is mandatory for any valid inference about activation locations within the anatomical reference frame. Most of the functional MR imaging studies employ EPI sequences, which are subject to geometrical distortions due to susceptibility effects near the tissue-air and tissue-bone interfaces (Jezzard and Balaban, 1995). This prevents the reduction of the co-registration of the distorted functional to an undistorted structural image to a 6-parameter rigid-body problem. The pattern of geometric distortions in

an EPI sequence depends not only on the anatomy of the subject, but also on field strength B_0 and the measured matrix size. The distortions measured with a 3 Tesla scanner and a 128×128 -matrix would be greater than in images measured on a 1.5 Tesla scanner with a 64×64 -matrix. A principled solution to the distortion problem is to apply an unwarping scheme to correct for the geometrical distortion. This can be done by acquisition of a field map using a double-echo gradient recalled echo image (Jezzard and Balaban, 1995) so that the field inhomogeneity map can be calculated. This field inhomogeneity map is then employed to unwarp the geometrical distortions in the phase encoding direction, assuming that the distortions in read gradient direction are negligible.

3.5 Applications

Surface reconstruction is a computationally intensive and partially interactive process. Usually, the reconstructed surface of one hemisphere consists of roughly 150,000 vertices and 320,000 faces, which requires a system with excellent computing power and graphics performance for generation and visualization (at least in December 2000). Given such a system, the interactive elimination of defects can be time-consuming. A legitimate question is then, what the benefits of this effort are. These benefits are now discussed.

Generally, the approximation of the cortical surface is a representation of the structure and topology of an individual's grey matter sheet. This information is implicitly accessed by the human observer of a structural MR image, but this anatomically informed knowledge cannot be specified in the original voxel-space. Surface reconstruction provides an explicit description of the topological structure of an individual's grey matter sheet and in particular, this representation is accessible to further automated manipulation.

The computation of geodesic (minimal) distances on the reconstructed surface is a non-trivial problem. As a first approximation, one can estimate the minimal distance along the edges of the faces using Dijkstra's shortest-path algorithm (Fischl et al., 1999b). Obviously, this is an overestimate of the true shortest path. A better method is proposed in (Kimmel and Sethian, 1998), which can be used to find minimal paths along the reconstructed surface, i.e. the path does not need to follow edges, but can cross faces.

Given coregistered statistical maps in (functional) voxel-space, one can attempt to project these onto the individual's surface or its transforms as in several studies, mostly concerning the visual cortex (Serenio et al., 1995; Engel et al., 1997; Tootell et al., 1997). The assessment of retinotopy in voxel-space is a demanding task, but simplified in unfolded vertex-space. Other examples, where this kind of projection provides a useful tool, include the characterization of somatotopical and tonotopical maps.

However, the projection of statistical maps onto surfaces is a complex issue. A simple scheme would involve interpolating the intensity value of the surrounding statistical voxel intensities at each vertex. This is comparable to overlaying the statistical map onto an individual's structural MRI. An important point to note here is that

the smoothness of the statistical map, which is primarily due to the functional resolution and any exogenous lowpass filter, can cause an incorrect topological mapping and spoils the benefit of the reconstructed high-resolution surface. There are several countermeasures to tackle this problem: (i) one can increase the applied threshold of the statistical map to generate more focal projections, but this can exclude weaker (but still significant effects) from the characterization, (ii) one can decrease the width of the applied lowpass filter or even omit it, but this means that sensitivity for extended activations is decreased and (iii) one can choose a projection surface different from the grey matter surface, e.g. the white-grey interface such that large activation clusters originating in the grey matter or in adjacent CSF are projected onto a smaller and possibly more interpretable set of vertices. However, the use of all these techniques speaks to the fundamental underlying issue that a projection from three-dimensional functional voxel-space to a high-resolution surface does not change the analysis' sensitivity or its intrinsic spatial resolution. To conclude, the projection of statistical data in voxel-space onto surfaces can be a helpful visualization tool for functional MRI studies, if the intrinsic spatial structure of the activations is rendered more obvious on the surface. However, the primary features of such a surface, i.e. its sub-voxel resolution, intrinsic topological structure and geometry are not used during the statistical analysis per se.

The projection of a brain surface to a plane offers the advantage of a surface coordinate system such that distance and orientation can be easily assessed on the flat map. The unfolded surface has been proposed as a reference system or atlas system (van Essen and Drury, 1997), because the coordinate system of a flat map directly reflects the topology of the grey matter sheet. But as discussed by Fischl et al. (1999b), the price paid is high: (i) The cuts introduce a topologically (and geometrically) incorrect mapping, i.e. cuts destroy neighbourships, and (ii) the flattened cortical surface is not convex, i.e. there are points in the coordinate systems that have no reference on the folded surface.

Another coordinate system, employing the unfolded map, is given by the projection of the surface onto a sphere or ellipsoid (Fischl et al., 1999b). Here, no cuts prior to the projection are necessary and the topology can be fully retained. However, the impact of residual distortions is still an issue. Distance and orientation can be assessed, even though not as conveniently done on a plane.

Once the reconstructed surface is unfolded, a surface-based normalization can be applied to transform the individual's map to the reference system. A disadvantage of any surface-based reference systems is that the distribution of distortions are dependent on the definition of the atlas, i.e. it might be that distortions due to the unfolding are large in brain areas one is particularly interested in.

Another useful application of surface-based normalization is morphometry analyses based on surface features like curvature or thickness of the grey matter sheet (Thompson et al., 1997).

In (Fischl et al., 1999a), it has been proposed that surface-based normalization can improve the sensitivity of functional analyses in comparison to affine normalization in voxel-space. Although it would have been more appropriate to compare a (non-linear) surface-based normalization to a non-linear normalization in voxel-space, this line of

research addresses the interesting question, whether a surface-based normalization provides a better normalization than one in voxel-space. This would be certainly the case, if functional areas are rigidly located in any subject to macroscopic landmarks given by the gyral pattern. Cytoarchitectonical studies (Amunts et al., 2000) show that this is e.g. approximately true for primary visual areas, but most certainly not in higher order areas (Amunts et al., 1999). This finding means that increasing the accuracy of the normalization with respect to aligning structural landmarks visible on the T_1 -weighted MRI is not necessary for multi-subject studies, because the variability of the exact location of functional areas over subjects with respect to structural landmarks is too high. This issue of anatomical variability is related to the results presented in this thesis and will be further discussed in chapter 7.

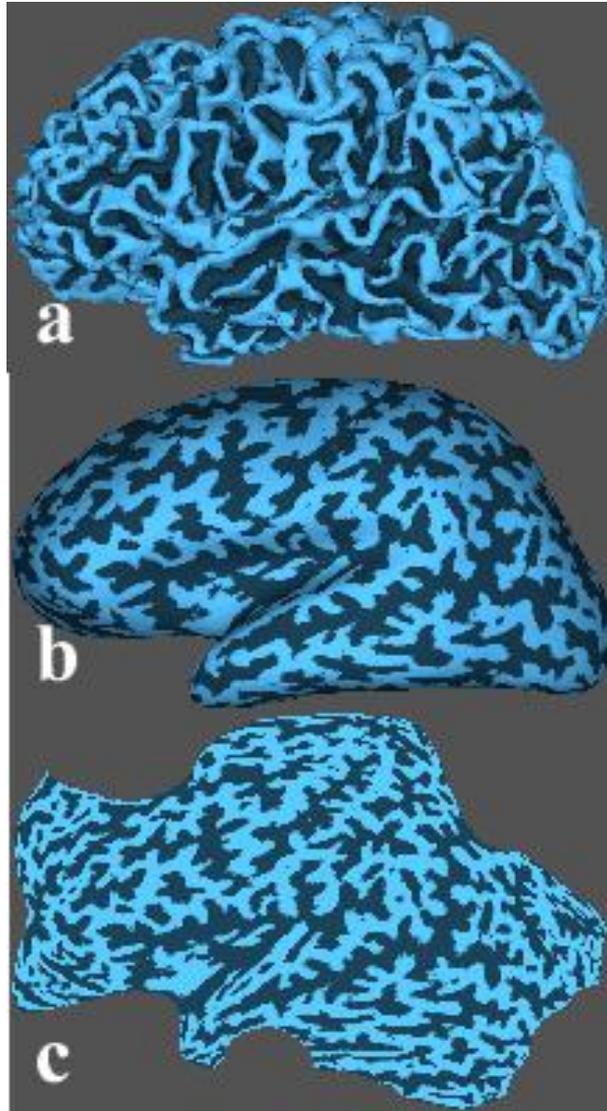


Figure 3.1: Different surfaces representations of one hemisphere. (a) folded surface S_G , (b) inflated surface S_I and (c) flattened surface S_F .

Chapter 4

Model specification

4.1 Introduction

The problem addressed by anatomically informed basis functions can be stated as follows:

Given the assumptions that: (i) the measured metabolic effects are located somewhere on the cortical surface and (ii) can be characterized by a spatially smooth distribution on that surface, what is the spatial distribution of the metabolic signal, which best explains the functional observations?

If both assumptions are valid, one can describe an anatomically informed model to characterize any measured fMRI or PET data directly on the cortical surface. As it is shown later there are three major advantages of the technique when compared to voxel-based methods: (i) the method does not use an isotropic spatial lowpass filter, but defines a spatially variant model at each voxel-position, (ii) the method can enhance the resolution of statistical results with respect to the underlying neuroanatomy and (iii) AIBF can be used to increase the sensitivity of functional data analysis. This includes both single-subject and multi-subject analyses.

Additionally, AIBF can also be used to implement an anatomically constrained least-squares deconvolution for low-resolution data. This applies when the intrinsic point spread function (PSF) is large relative to voxel size, which is e.g. the case for PET data.

In the following, the approach is described in three sections. First, the spatial model that embodies anatomical a priori knowledge about a functional observation is specified. In the second section, the anatomical model is fitted to a series of functional volumes observed over time, to generate a spatiotemporal parameter matrix. The third section deals with inference about effects in the time domain. In this work, it is assumed that the spatiotemporal basis set can be factorized into a spatial and temporal basis set, where the latter corresponds to conventional regressors in a design matrix. In this way, one can first estimate the spatial distribution of sources at each time point and then

estimate, and make inferences about, their expression over time.

4.2 The spatial domain

4.2.1 The spatial equation system

It is assumed that the individual grey matter surface $S_G = (V_G, F)$ is a good approximation to the real underlying grey matter surface in the individual's anatomical space. Given a functional observation vector Y in voxel-space, the overall aim is to estimate a smooth distribution f on the vertices V_G of the reconstructed grey matter surface that best (in some sense) explains the observed Y , i.e.

$$Y = g(f(V_G)) + \epsilon \quad (4.1)$$

where ϵ is the residual component and g is an operator (to be described later), which transforms intensities defined in vertex-space to the voxel-space of Y . Typically, one reconstructed hemisphere consists of roughly 130,000 vertices, whereas the number of observed grey matter voxels in one hemisphere (e.g. 20 slices, 128×128 -matrices, voxel size $1.8 \times 1.8 \times 3 \text{ mm}^3$) is much less, roughly 12,000, i.e. the problem is underdetermined. One key to solve this problem is the assumption about the smoothness of f , i.e. any focal neuronal activation will manifest as a smooth distribution on or close to the cortical surface. This smooth distribution can be modelled by less parameters N_p than the number of vertices N_{V_G} so the idea is to approximate f by a model that is determined by a small set of parameters.

In this thesis, the distribution f was modelled with two kinds of spatial bases, both capable of modelling smooth distributions. The first was a linear combination of local, smooth and partially overlapping spatial basis functions defined on V_G , which are called *local* basis functions in the context of this thesis. Fig. 4.1 illustrates the concept behind this approach. The other model consisted of a set of discrete cosine functions at various low frequencies, which are named *global* basis functions.

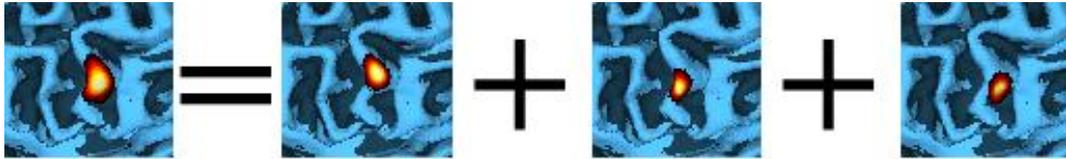


Figure 4.1: Concept of smooth overlapping local spatial basis functions defined on the surface S_G : The distribution on the left hand side can be expressed as a linear combination of the three basis functions on the right.

A necessary prerequisite for the specification of spatial basis functions on the surface S_G is the existence of a two-dimensional coordinate system, but due to the folded nature of the grey matter surface, there is no simple way to define such a coordinate system.

However, the cortical flat map $S_F = (V_F, F)$ is the result of a coordinate transformation from the vertex coordinates V_G to a plane, i.e. for any vertex coordinate v_G on S_G its associated location on S_F is stored in v_F . Since there is a two-dimensional coordinate system available on the cortical flat map S_F , we define the basis functions on S_F and re-project them onto the cortical sheet S_G by a coordinate exchange.

4.2.2 Specification of basis functions on the flattened cortex

Three different sets of basis functions were used. Two of them were sets of local basis functions and the third set consisted of global basis functions. In the following, the basis function sets are defined and their features are outlined.

There are three conditions a local basis function set must conform to: (i) each basis function must cover a *small patch* of S_F , (ii) the linear combination of basis functions is able to model smooth distributions on S_F and (iii) the approximation of f at any vertex location is implemented as sparsely as possible. Any form of two-dimensional Gaussian, sinc-function or similar lowpass-kernel functions fulfill the first two criteria, provided that each basis function overlaps with its neighbours to a certain degree. The third condition means that this overlap is not unnecessarily large.

Two different sets of local basis functions were used. One employed circular Gaussian basis functions. The other set consisted of circular Sinc-functions. Both sets were arranged in a hexagonal configuration such that the centre of each basis function had a fixed (user-specified) distance to the centre of its six neighbours.

The global set of basis functions was a discrete cosine set. Each basis function covers the whole surface of S_F and encodes two specific spatial frequencies, one in the x-direction and one in the y-direction. The only parameter the user has to choose, when using this global basis function set, is the range of spatial frequencies the model contains.

Gaussian basis functions

Given a two-dimensional coordinate system on the cortical flat map S_F , the basis function b_F^j with its centre at coordinates (x_j, y_j) is defined by

$$b_F^j(x, y) = c_1 \exp\left(\frac{-((x - x_j)^2 + (y - y_j)^2)}{2w^2}\right) \quad (4.2)$$

where c_1 is a constant and w is the (user-specified) width of the Gaussian basis function in the x- and y-directions. The hexagonal pattern of the centres of the basis functions are defined by induction, i.e. given that (x_j, y_j) is the centre of b_F^j , then the centres of its six neighbouring basis functions are $(x + d/2, y + d_o)$, $(x + d, y)$, $(x + d/2, y - d_o)$, $(x - d/2, y - d_o)$, $(x - d, y)$, $(x - d/2, y + d_o)$, where d is the fixed distance between centres and $d_o = \sin(60^\circ)d$. The basis functions are discretely sampled at vertex positions.

The position of the first basis function is chosen arbitrarily. A basis function is defined only if its centre is enclosed by a face of the cortical flat map S_F and its actual support

area on S_F is larger than some specified proportion of its potential support area. The latter condition can be used to implement boundary conditions.

Sinc basis functions

Similarly, a Sinc basis function with its centre at (x_j, y_j) is defined by

$$b_F^j(x, y) = \text{sinc}(x - x_j, w) \cdot \text{sinc}(y - y_j, w) \quad (4.3)$$

where

$$\text{sinc}(x, w) = \begin{cases} 1 & x = 0, \\ \frac{\sin(x/w)}{x} & \text{otherwise.} \end{cases} \quad (4.4)$$

The width w is specified by the user and the locations of the centres are specified as in the case of Gaussian basis functions.

Specification of discrete cosine basis functions

Each basis function of the global set encompasses the whole surface S_F . The two dimensional discrete cosine set ([Gonzalez and Wintz, 1987](#)) is defined by

$$dct(x, y, u, v) = \begin{cases} \frac{1}{N_{max}} & u = 0, v = 0, \\ \frac{1}{2N_{max}^3} \cos((2x + 1)u\pi) \cos((2y + 1)v\pi) & \text{otherwise} \end{cases} \quad (4.5)$$

where x and y specify coordinates of S_F , $u = 0, \dots, N_u$ and $v = 0, \dots, N_v$, N_{max} is the maximum possible order of the discrete cosine set, N_u and N_v are the maximum order basis functions incorporated into the set in the x- and y-directions. N_{max} is determined by the maximum side length of S_F measured in mm on the cortical surface. Additionally, $0 < N_u, N_v < N_{max} - 1$. In practice, because one assumes that the distribution of spatial frequencies does not differ in the arbitrarily chosen x- and y-direction, the user chooses an order of the DCT-set (discrete cosine transform) that is interpreted as the order for the direction (in the x- or the y-direction), in which the flattened surface has maximum length. The order of the other direction is adjusted accordingly such that the maximum frequencies in the x- and y-directions are roughly equal. Note that the actual spatial frequencies with respect to S_F modelled by a basis set is not only a function of N_u and N_v , but also of the actual size of S_F in the x- and y-direction.

4.2.3 Reprojection and transformation

After specification of the basis functions on the cortical flat map $b_F^j, j = 1, \dots, N_p$, where N_p is the number of basis functions, each b_F^j is re-projected onto the original folded surface S_G by a vertex coordinate exchange between V_F and V_G to give folded basis functions b_G^j . Since the b_G^j are still in vertex-space, the next step is to embed each

basis function b_G^j into the voxel-space. To do this, an operator g (Eq. 4.1) is defined, which integrates over each voxel k the surface of a given basis function b_G^j multiplied by the height of the basis function, i.e. g returns the integral of each basis function in folded vertex-space encompassed by the voxel.

To do this, one assumes that the activity distribution on a given face is a linear function of activity in the three vertices of this face. Then an $N_K \times N_{V_G}$ matrix G can be defined, where N_K is the number of functional voxels. Each column j of G specifies the observed activity in voxel-space given some activity in vertex j .

The matrix-vector multiplication Gb_G^j transforms the basis functions b_G^j in vertex-space to b_Y^j ($b_Y^j = b_{Y_1}^j, \dots, b_{Y_{N_K}}^j$) in voxel-space.

The process of transforming a Gaussian basis function from its initial representation on the cortical flat map, to its intermediate state on the folded cortical map, and to its final representation in fMRI voxel-space is shown in Fig. 4.2.

4.2.4 Convolution by an additional point spread function

If the width of the image point spread function (PSF) is larger than the voxel size, one has to apply an additional convolution L_I to the basis functions in voxel-space b_Y^j , giving $b_{L_I}^j$. This is not necessary for fMRI, since here it is assumed that $L_I = I$, where I is the identity matrix such that $b_Y^j = b_{L_I}^j$.

For example with PET (positron emission tomography) data, one could model the positron travel, the effects of the measurement process and the reconstruction filter by convolution of b_Y^j with an isotropic lowpass filter kernel. Obviously, a spatially stationary kernel L_I is only an approximation to the underlying presumably non-stationary PSF, but the important point is that the width of L_I in voxel-space roughly matches the image smoothness induced by the unknown image PSF. In PET one generally assumes that the PSF is stationary.

This basis set is now used to model, in a linear combination, the functional observation Y .

4.2.5 Estimating the spatial distribution of activity for one time point

Let A_G be a $N_{V_G} \times N_p$ -matrix, where column j of A_G is the basis function b_G^j in vertex-space, i.e. $A_G = [b_G^1, \dots, b_G^{N_p}]$.

As described above, let G be an $N_K \times N_{V_G}$ -matrix, where G encodes the effect of activity in vertex-space observed in voxel-space. Let L_I be an $N_K \times N_K$ convolution matrix to model a stationary image PSF. Let L_E be a $N_K \times N_K$ convolution matrix to model an additional exogenous smoothing filter applied to both the data and model. $L = L_E L_I$ implements the effective PSF after spatial smoothing. $A_L = n(LG A_G)$, where n normalizes each column of a matrix such that the sum of squares of the

column is unity.

Then, the model is

$$L_E Y = A_L \beta + L_E \epsilon \quad (4.6)$$

where Y is the N_K -dimensional observation vector in voxel-space, β is a N_p -dimensional parameter vector and ϵ is a N_K -dimensional error vector.

A_L is then a $N_K \times N_p$ -matrix, where column j of A_L is the convolved basis function $n(b_L^j)$ in voxel-space. Similarly, let A be a $N_K \times N_p$ -matrix, where column j of A is the basis function $n(b_Y^j)$. For fMRI, it is assumed that $L_I = I$.

The unknown parameter vector β is assessed by a regularized least-squares estimate (Press et al., 1992), which yields

$$\hat{\beta} = (A_L^T A_L + \lambda W^T W)^{-1} A_L^T L_E Y \quad (4.7)$$

where λ is a regularization factor and W is a $N_W \times N_p$ -matrix imposing additional smoothness constraints on the solution (see next section).

4.2.6 Regularization

Depending on the specified basis function matrix A_L , regularization can help to stabilize the solution and prevent wildly oscillating parameters over space. Most regularization techniques specify some smoothness constraint in W as prior knowledge about the underlying solution. When using basis functions as proposed in this thesis, this is not strictly necessary, because a smooth set of basis functions already implements a smoothness constraint. As a special case, however, towards the limit that A_G is an identity matrix, i.e. no basis functions are used, regularization of the solution becomes absolutely essential, because the solution would be otherwise underdetermined.

As it will be shown in sec. 5.2.4, regularization should be used when the overlap between local basis functions is high. In such cases, regularization of the least-squares solution with an additional smoothness constraint decreases the spatial variance of the estimated parameters $\hat{\beta}$ such that the influence of high spatial frequencies within the spectrum of the basis functions is diminished.

Two regularization constraints have been used. The first was a zeroth order regularization with respect to the parameter space. Here, $W = I_{N_p}$, thereby implicitly imposing the constraint on the solution to minimize the 2-norm of the vector $\hat{\beta}$. The second regularization was of first order in the voxel-space, i.e. minimizing the sum of the 2-norm of the first partial derivatives of the fitted image in all three directions. This can be implemented by setting $W = \sum_{i=1}^3 (\dot{A}_L)_i$, where i denotes the direction of the derivative.

Given the regularization constraints, a question remains how λ should be chosen. There is a literature describing schemes to find the *optimal* λ given that Y , A_L and W are known (Harville, 1974; Hansen, 1992; Engl and Grever, 1994; Press et al., 1992).

One computationally inexpensive approach (Press et al., 1992) is to estimate λ by

$$\lambda = \frac{\text{trace}(A_L^T A_L)}{\text{trace}(W^T W)} \quad (4.8)$$

A more sophisticated algorithm is given by the restricted maximum likelihood (ReML) algorithm (Harville, 1977).

The ReML method uses an iterative optimization as follows: Initialize the estimate of two unknown hyper-parameters θ_1 and θ_2 with two positive values and estimate β by

$$\hat{\beta} = (\theta_1^{-1} A_L^T A_L + \theta_2^{-1} W^T W)^{-1} \theta_1^{-1} A_L^T L_E Y \quad (4.9)$$

Compute the residuals

$$\mathbf{r} = \begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{pmatrix} = \begin{pmatrix} Y - A_L \hat{\beta} \\ W \hat{\beta} \end{pmatrix} \quad (4.10)$$

The degrees of freedom for the estimate of β are given by

$$\nu_1 = \text{trace}((\theta_1^{-1} A_L^T A_L + \theta_2^{-1} W^T W)^{-1} \theta_1^{-1} A_L^T A_L) \quad (4.11)$$

$$\nu_2 = N_p - \nu_1 \quad (4.12)$$

The updates of the estimates of θ_1 and θ_2 are

$$\hat{\theta}_1 = \frac{\mathbf{r}_1^T \mathbf{r}_1}{N_K - \nu_1} \quad (4.13)$$

$$\hat{\theta}_2 = \frac{\mathbf{r}_2^T \mathbf{r}_2}{N_W - \nu_2} \quad (4.14)$$

After convergence, λ (Eq. 4.7) is assessed by

$$\lambda = \left(\frac{\hat{\theta}_1}{\hat{\theta}_2} \right)^{1/2} \quad (4.15)$$

4.2.7 Global basis function

As an extension to Eq. 4.6, one can model a globally coherent component by adding a global basis function to the basis set. In this thesis, two different global basis functions have been evaluated. Both are an extension to Eq. 4.6. The first global basis function is uniform over the grey matter encompassed, i.e.

$$A_L = n(LG[A_G|\mathbf{1}]) \quad (4.16)$$

and the second uses the mean of the observed images as a basis function such that

$$A_L = n(L[G A_G|\bar{Y}]) \quad (4.17)$$

The regularization matrix W is modified accordingly by having zero entries for the global basis function, i.e. the fit of the global basis function is not regularized or penalized.

4.3 The temporal domain

Given a series of functional observations Y_1, \dots, Y_{N_Y} , we estimate for each Y_i a parameter vector $\hat{\beta}_i$ by Eq. 4.7 to assemble an $(N_Y \times N_p)$ estimated parameter matrix

$$B^T = [\hat{\beta}_1 | \dots | \hat{\beta}_{N_Y}] = (A_L^T A_L + \lambda W^T W)^{-1} A_L^T L_E [Y_1 | \dots | Y_{N_Y}] \quad (4.18)$$

which represents the estimate of functional observations, projected into the space of anatomically informed basis functions.

4.3.1 Reprojection of parameters and deconvolution

The estimated signal in voxel-space subject to convolution by L is given by

$$B_L = B A_L^T \quad (4.19)$$

Similarly, if $L_E \neq I_{N_K}$, one can project the estimated parameters to voxel-space subject to convolution with L_I (only) by

$$B_{L_I} = B A_{L_I}^T \quad (4.20)$$

The estimated signal in voxel-space, but without convolution by L is

$$B_A = B A^T \quad (4.21)$$

By omitting the spatial convolution matrices L_E or L_I one is effectively implementing a least-squares deconvolution in the sense that the estimates are those that obtain prior to convolution.

Equivalently, let the columns of matrix A_{vertex} be the folded basis functions in vertex-space b_G^j , then the estimated signal in vertex-space is given by

$$B_{vertex} = B A_{vertex}^T \quad (4.22)$$

Eqs. 4.19 to 4.22 can also be expressed in an alternative way. For example, by substituting B in Eq. 4.19,

$$B_L = [Y_1 | \dots | Y_{N_Y}]^T L_E^T A_L (A_L^T A_L + \lambda W^T W)^{-T} A_L^T \quad (4.23)$$

Let

$$P_{AL} = L_E^T A_L (A_L^T A_L + \lambda W^T W)^{-T} A_L^T \quad (4.24)$$

such that

$$B_L = [Y_1 | \dots | Y_{N_Y}]^T P_{AL} \quad (4.25)$$

P_{AL} is a $N_K \times N_K$ matrix that projects the data convolved with L_E to the fitted data B_L . Note that P_{AL} represents an anatomically informed spatially variable filter kernel at each voxel, where each column k of P_{AL} encodes the spatial filter kernel image at voxel k . A projection matrix that represents an anatomically informed deconvolution of the data (Eq. 4.21) is given by

$$P_A = L_E^T A_L (A_L^T A_L + \lambda W^T W)^{-T} A^T \quad (4.26)$$

4.4 Inferences about evoked responses over time

To make inferences about evoked responses, one has to characterize the parameter matrix B in the temporal domain. This could be done at four levels: (i) the first is in the parameter space (Eq. 4.18), (ii) the second is in the vertex-space (Eq. 4.22) and (iii) one can also project B back into voxel-space given by Eq. 4.21 or (iv) into the voxel-space specified by Eq. 4.19 and proceed with a conventional voxel-based analysis in both cases.

Either a univariate or a multivariate technique could be used at all four levels.

In the following, the different levels will be called parameter space (Eq. 4.18), vertex-space (Eq. 4.22), A -space (Eq. 4.21) and A_L -space (Eq. 4.19). The latter two spaces are called voxel-spaces.

4.4.1 Univariate inferences

Given a temporal design matrix M , one can model the spatial parameter matrix over time by

$$B = M\gamma + \epsilon_B \quad (4.27)$$

such that the least-squares estimate of γ is given by

$$\hat{\gamma} = (M^T M)^{-1} M^T B \quad (4.28)$$

This estimate can then be used to make inferences by using the framework of voxel-based methods as described in sec. 2.2.3.

Interestingly, any other inference in the spaces defined by A_L , A and A_{vertex} could be directly derived from $\hat{\gamma}$, because e.g. for the space defined by A the projected estimated parameter matrix B_A (Eq. 4.21) can be modelled by

$$B_A = M\gamma_{B_A} + \epsilon_{B_A} \quad (4.29)$$

such that using Eq. 4.27

$$(M\gamma + \epsilon_B)A^T = M\gamma_{B_A} + \epsilon_{B_A}$$

which means

$$\gamma A^T = \gamma_{B_A} \quad (4.30)$$

$$\epsilon_B A^T = \epsilon_{B_A} \quad (4.31)$$

In short γ needs to be estimated only once and the other temporal parameter matrices γ_{vertex} , γ_A and γ_{A_L} in the spaces of A_{vertex} , A and A_L can be derived by a simple projection. Building on this, any other variables (e.g. the residuals) necessary for the assessment of a t- or F-field in the vertex- or voxel-spaces can be assessed by projection of the associated variable in the parameter space defined by B . This procedure can only be employed, if the temporal design matrix M is the same for projection and parameter space. The other straight-forward way of assessing a t- or F-field in vertex- or voxel-space is to project B into voxel- or vertex-space (e.g. Eq. 4.21) and estimate the temporal parameters in this space. The advantage of projecting $\hat{\gamma}$ instead of estimating e.g. γ_A based on B_A is that one avoids reiterated estimation of the parameters, which includes handling, multiplication and inversion of potentially large temporal design matrices as in the case of event-related fMRI.

The problem of multiple comparisons over voxels can be overcome by application of the results from the theory of random fields as is standard practice in SPM. The smoothness of non stationary residuals in voxel-space B_A and B_L can be estimated using results from (Worsley et al., 1999). Using anatomically informed basis functions, one has to take into account that B_A features a spatially non-stationary smoothness distribution because of the re-folding of the flat basis functions b_j^F into voxel-space. In (Worsley et al., 1999), it is shown that it is sufficient to estimate the smoothness based on the expectation of the determinants of the partial derivatives of the residual fields to estimate the (local) smoothness at any voxel. This technique allows the appropriate adjustments of the p-values in a statistical image with non-stationary smoothness using results from differential topology.

The theory of random fields cannot be applied to the graph of t- or F-values based on the parameter matrix $\hat{\gamma}$, because this representation is not a discretized version of an underlying continuous field. The *spatial* noise covariance matrix (estimated on the standardized residual components) is related to the smoothness of a statistical field, but only the projection to some image space establishes the assumption of a sampled version of a continuous field. The Bonferroni correction can be applied, although it is

clear that this method yields overly conservative p-values for a single basis function due to the large covariance between spatially neighbouring parameters of (e.g. Gaussian) basis functions. Additionally, an important question is whether a significant p-value of a basis function supports the interpretation of the functional data. In the case of local basis functions, the interpretation of a significant effect in parameter space can be that there is some activation in the volume defined by the basis function. In other words, the localization property of a single Gaussian or Sinc basis function is poor as compared to a test within a voxel. Discrete cosine basis functions have no localization power at all. The information given by the parameter of a single DCT basis function is that there is some spatial frequency component in the direction of the cortical sheet, which on its own unlikely to be useful. These considerations speak to inferences about all the parameters collectively, i.e. the inference that the brain has responded *somewhere*. To adopt this approach one leaves the mass-univariate methodology of SPM and turns to multivariate inferences.

4.4.2 Multivariate inferences

In this section, an approach making multivariate inferences about the parameter matrix B is described. This multivariate method was first described in (Friston et al., 1995b) and is based on a combined approach of singular value decomposition (SVD), multivariate analysis of covariance (ManCova) and canonical variates analysis (CVA) (Chatfield and Collins, 1980).

As with a conventional analysis in voxel-space (e.g. in SPM99), various types of temporal basis functions or regressors were used: (i) condition encoding vectors, (ii) a constant and (iii) a vector of global mean intensities. The condition encoding basis functions are the covariates of interest, the remainder generally represent covariates of no interest. If the data were acquired with fMRI, a temporal band-pass filter is applied to the functional data prior to fitting the model.

In what follows it is assumed that the band-pass filter accounts for temporal auto-correlations, which are especially prominent at lower frequencies (Zarahn et al., 1997; Aguirre et al., 1997).

Let the temporal basis functions be $M = [M_1 M_2]$, where M is an $N_Y \times N_M$ -matrix, N_M is the number of temporal basis functions, M_1 and M_2 are $N_Y \times N_{M_1}$ - and $N_Y \times N_{M_2}$ -matrices, M_1 and M_2 are orthogonal to each other, and $N_M = N_{M_1} + N_{M_2}$. M_1 contains the covariates of interest and M_2 the covariates of no interest. A multivariate analysis proceeds as follows:

First, the effects of no interest are removed by

$$B_C = B - M_2(M_2^T M_2)^{-1} M_2^T B \quad (4.32)$$

The SVD of the *corrected* parameter matrix B_C is given by

$$B_C = B_U B_V B_W^T \quad (4.33)$$

where B_U and B_W are orthogonal matrices and B_V is a diagonal matrix. Let $B_P = B_U^J B_V^J$ a $N_Y \times N_J$ -matrix, where the columns of B_P contain the temporal expression of the first N_J spatial modes over observations. N_J is found by thresholding the associated singular values contained in B_V as described by Friston et al. (1995b) and B_U^J and B_V^J are the accordingly reduced versions of B_U and B_V . In neuroimaging this dimension reduction is necessary to ensure that the number of observations greatly exceeds the number of variates.

Then the temporal model is given by

$$B_P = M_1 \gamma + \epsilon_M \quad (4.34)$$

where γ is a $N_{M_1} \times N_J$ -matrix and ϵ_M a $N_Y \times N_J$ -error matrix.

The least-squares estimate of γ is

$$\hat{\gamma} = (M_1^T M_1)^{-1} M_1^T B_P \quad (4.35)$$

Then the sums of squares and products due to error are

$$R_E = (B_P - M_1 \hat{\gamma})^T (B_P - M_1 \hat{\gamma}) \quad (4.36)$$

The sum of squares and products due to the effects of interest are

$$H = (M_1 \hat{\gamma})^T (M_1 \hat{\gamma}) \quad (4.37)$$

and the sum of squares and products under the null hypothesis (that the effects due to M_1 do not exist) are

$$R_0 = B_P^T B_P \quad (4.38)$$

The significance of the effects of interest can be tested with Wilk's Lambda

$$\Lambda = \frac{|R_E|}{|R_0|} \quad (4.39)$$

as described in (Chatfield and Collins, 1980).

The characterization of the significant effects employs CVA, i.e. one finds a matrix of canonical images $Q = [\mathbf{q}_1, \dots, \mathbf{q}_{N_J}]$ such that the variance ratio $(\mathbf{q}_m^T H \mathbf{q}_m) / (\mathbf{q}_m^T R_E \mathbf{q}_m)$ is maximized successively for $m = 1, \dots, N_J$ under the condition that $cov(\mathbf{q}_m, \mathbf{q}_n) = 0$ for any m, n with $1 \leq m, n \leq N_J$ and $m \neq n$.

4.5 Summary

The outline of the proposed algorithm is presented in Figs. 4.3 and 4.4 in the form of flow charts. In Fig. 4.3, the information flow from the measured functional and structural images to the temporal analysis is described. Fig. 4.4 represents the critical

part of AIBF model generation in more detail. Both flow charts are described in the following.

In Fig. 4.3, the processing graph can be divided into three parts. On the left hand side, the aligned functional images are preprocessed, while on the right hand side the structural image is transformed to the surface representation and the AIBF model is generated. The fusion of both information streams is implemented in the model fitting stage (see Eq. 4.6) such that images can be projected to inference space (Eqs. 4.19 to 4.22) and a temporal analysis can be performed. More specifically, as a first step, the mean of the functional (aligned) images is computed and co-registered to the structural image. Optionally, this mean image can be anatomically normalized. The normalization transformation is then applied to all functional images. On the structural side, the brain surface is reconstructed from the T_1 -weighted MR image. Optionally, the surface can be normalized using the normalization transformation derived from the functional mean image. The flattened representation is computed from the extracted (non-normalized) folded surface and the AIBF model is generated using the flattened surface, the normalized or non-normalized surface and the user specified AIBF parameters. This step is described in more detail in the text below (see also Fig. 4.4). After the model generation step, one has (normalized) functional images and a (normalized) AIBF model in the same voxel-space. Each image is fitted (Eq. 4.6) and projected to the desired space of inference (Eqs. 4.19 to 4.22). Inference can then be made using uni- or multivariate methods.

Fig. 4.4 depicts the processing flow of the AIBF model generation. A set of basis functions is specified on the flattened surface S_F . These two-dimensional basis functions are transformed (in vertex-space) to the folded surface representation and projected to voxel-space. Here, additional convolutions L_I and/or L_E can be applied to the basis functions. The convolved basis functions are finally assembled column-wise in a matrix and fitted to the functional data to give the estimated parameter matrix B (Eq. 4.18).

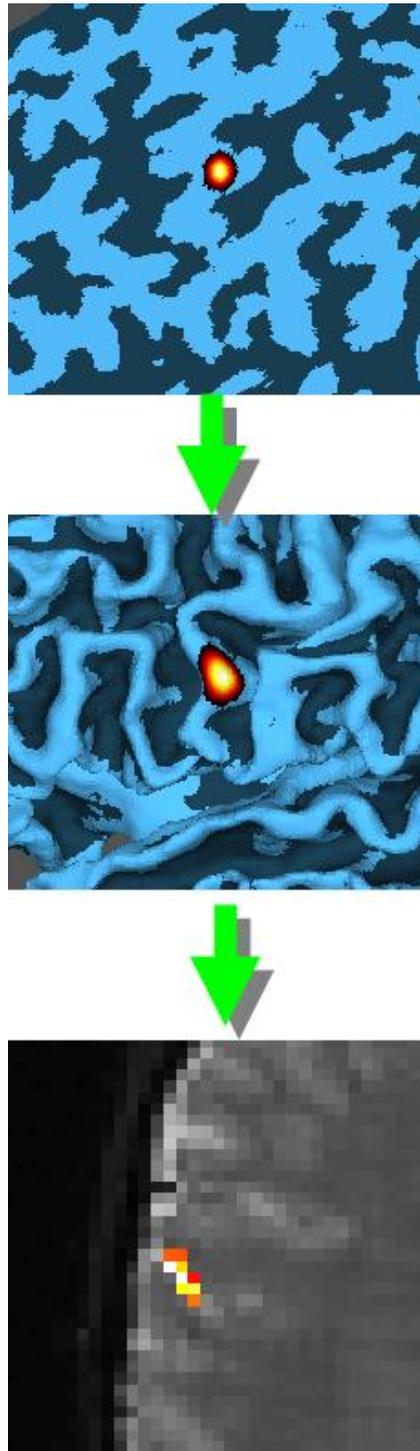


Figure 4.2: Different representations of a Gaussian basis function. (top) two-dimensional Gaussian kernel with 1 mm width in x- and y-direction, (middle) the same basis function transformed to the original reconstructed cortical surface, (bottom) a transverse slice of its representation in fMRI voxel-space.

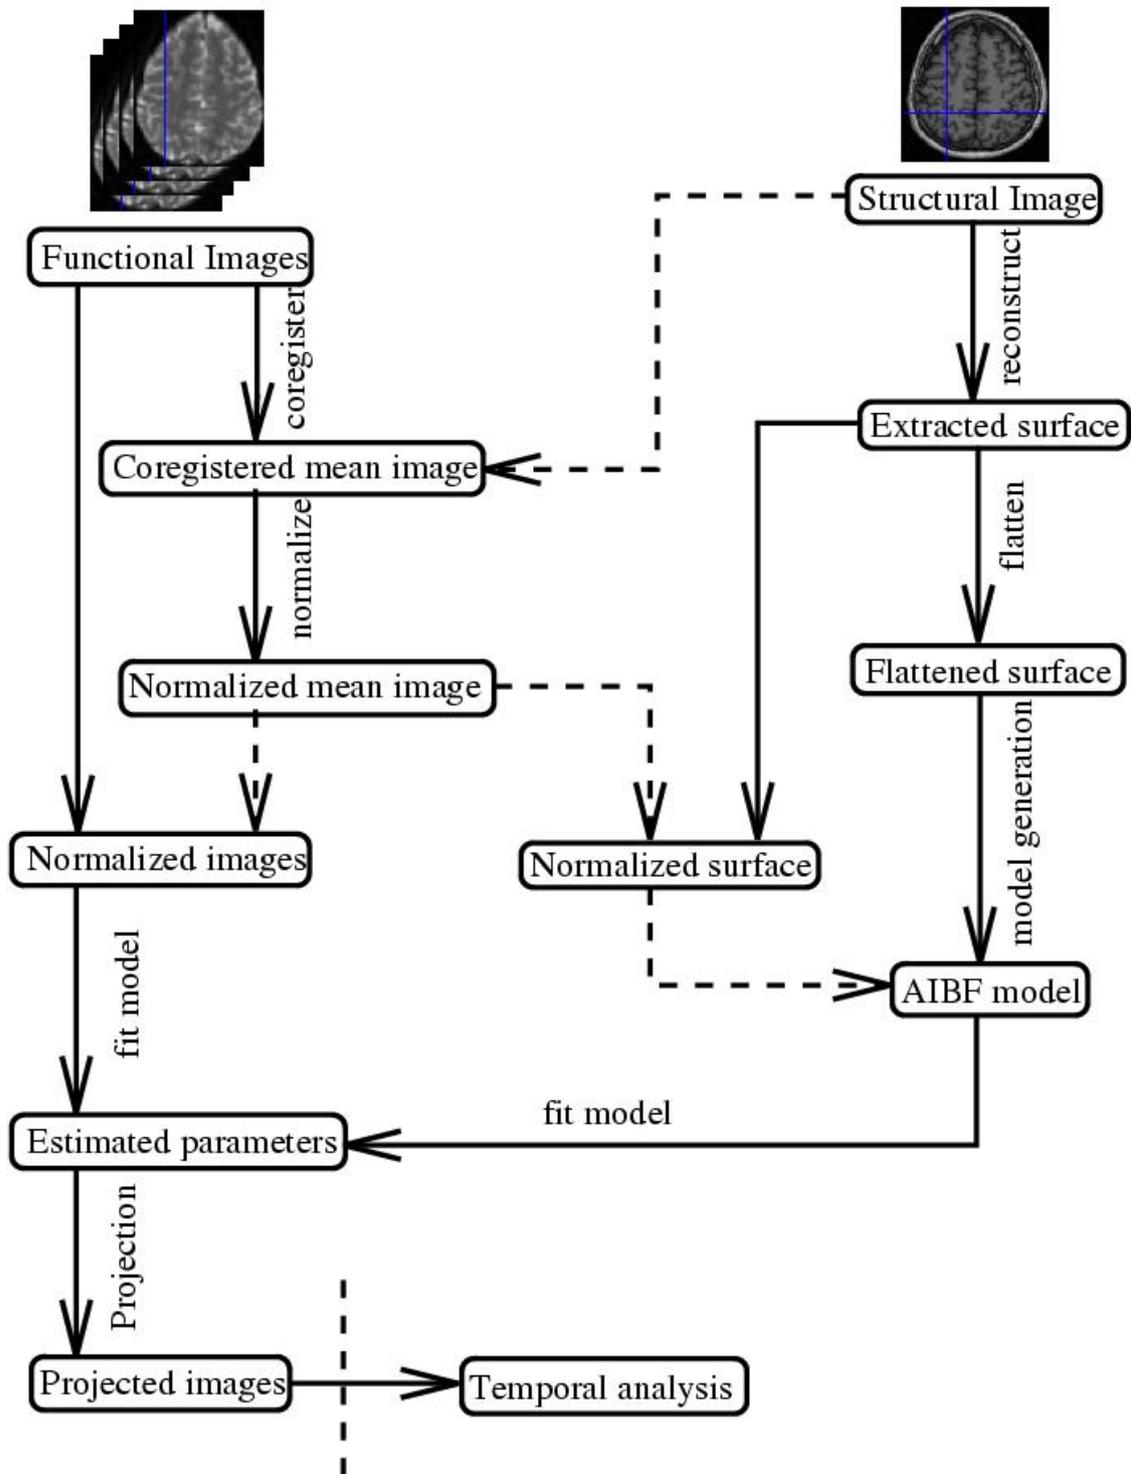


Figure 4.3: Flowchart describing the information flow from functional and structural images to the temporal analysis of the fitted and projected images. Solid arrows: transform of data, dashed arrows: transfer of information. A detailed description can be found in the text.

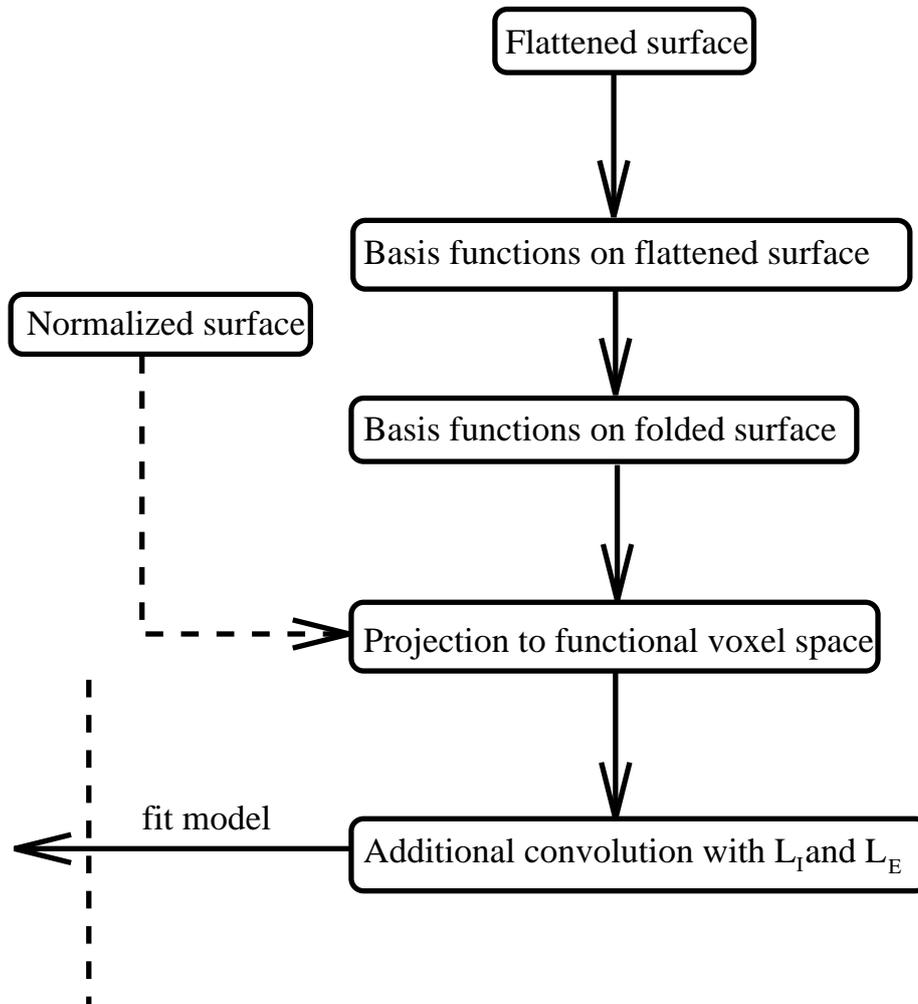


Figure 4.4: Flowchart describing the AIBF model generation. Solid arrows: transform of data, dashed arrows: transfer of information. A detailed description can be found in the text.

Chapter 5

Features of the method

In this chapter, various features of AIBF are described and demonstrated by applying it to some examples. In this context, using real data has the disadvantage that the true underlying signal is unknown such that it is difficult to make inferences about the validity of the method. Therefore, simulated data are used to examine the salient features, the technical behaviour, face validity and construct validity in relation to existing methods. Clearly, results with simulated data should be treated with the caveat that the simulation can only be an approximation to the real world.

In this chapter, the application to simulated data is used extensively to describe and discuss important features of the method in an ideal world, where the data largely conform to the assumptions required by the method. In the next chapter, AIBF is applied to real data.

5.1 Generation of simulated data

In this work, the process of generating functional data (fMRI and PET) is primarily based on one assumption: any activation is localized in or close to the grey matter sheet. The same assumption is made for the generation of the anatomically informed model, i.e. any potential activation elsewhere is not modelled and cannot be detected. Clearly, for fMRI data this assumption is likely to be incorrect, because it is known that some proportion of the signal comes from draining veins, which do not necessarily adhere to the grey matter sheet. This issue will be treated in the next chapter.

For both PET and fMRI, the information used to simulate an individual's functional data were a set of (null) functional data, a high-resolution structural T1-weighted image and the grey matter surface generated from this image. The following describes how the simulated data were generated for the fMRI and PET modalities.

5.1.1 Functional magnetic resonance data

Assuming that any activation arises from the grey matter sheet, one way of simulating functional data is to acquire a functional null (*activation-free*) data set and to add a simulated spatiotemporal signal based on regressors in the temporal design matrix and the reconstructed grey matter surface.

93 scans were acquired under a continuous rest condition, i.e. the subject was instructed to lie still and *think of nothing* for the whole session. The first 2 scans were discarded to avoid T_1 -effects, leaving 91 scans for the preprocessing and statistical analysis. As expected, this null data set showed no significant activations when analyzed with SPM99 based on a blocked design (maximum t-score 4.39, corresponding to a p-value of 0.93, corrected for multiple comparisons). In particular, there were no activations in the region of the left primary sensorimotor cortex, which was chosen as the region to introduce simulated activations.

The next step was the addition of some well-defined simulated activations. The left hemisphere of the individual's reconstructed cortical surface S_F and condition encoding temporal basis functions were used to define the activation signal. A subset of vertices of V_F (e.g. a circle of 3 mm diameter on the surface S_F) was selected and defined as the source of a BOLD effect. Then, a (signal) time series was specified in each vertex consisting of a compound of the condition encoding basis function. This signal was embedded into voxel-space as described in Chapter 4. The signal time-series in voxel-space was scaled such that the maximum intensity difference between conditions was $x\%$ of the intracortical global mean signal of the null data, where x was typically chosen from the range between 0 to 10%. Finally, the simulated signal in voxel-space was added to the null data.

5.1.2 Positron emission tomography data

The PET data was generated in a similar way. However, the author felt that the costs of scanning PET null data with respect to administering radioactively labelled water to a subject were too high compared to the benefit of generating a test vehicle for AIBF.

Alternatively, another appropriate procedure is to select an individual's PET scans from another activation study and assume that there is a region, which is virtually free of activation. Clearly, it cannot be shown that the activity in the selected ROI is not somehow affected by the experimental design, but at least a specific prior anatomical hypothesis and a statistical test can indicate that the effect size of an experimentally induced activation in the ROI is small.

Given such a PET data set, the same procedure was used as in the fMRI simulations. In addition, the signal time series were convolved with an isotropic Gaussian filter. This emulates the image point spread function generated by effects like positron travel, deviations from the 180° angle and the reconstruction filter. After the convolution, the simulated signal was scaled to match a given percentage of the grey matter mean. The FWHM of this Gaussian filter was chosen to be $[6\ 6\ 8]$ mm³ lying in the estimated

range of the effective PSF.

5.1.3 Ideal assumptions

Obviously, the simulated data resembles real data and has most of the properties of real data. However, there are several assumptions of the analysis that are met intrinsically by the generation process:

- the forward model of the BOLD-effect is correct, i.e. the effect of neuronal activations within the grey matter and the measurement process are modelled correctly by the matrix A_G
- All the signal is exclusively generated in the reconstructed grey matter surface
- The reconstruction of the grey matter surface is error free.
- Distortions between flat map and folded surface do not play a role.
- The realignment of the data is error free.
- fMRI: There are no distortions in the EPI sequences.
- The time course of the signal follows exactly the basis function used for the temporal analysis of the data.

5.2 Simulated studies

In the remainder of this chapter, five simulated experiments are presented. The first four use simulated fMRI data, whereas the fifth deals with simulated PET data. The first two show the performance of AIBF with respect to sensitivity and localization. In the third, the parameter space of different AIBF sets is explored. The fourth section discusses the effects of the regularization used to constrain the spatial solution as given in Eq. 4.7. The fifth section revisits some of the above issues using simulated PET data.

The analysis of all simulated data sets followed the same stages. All data sets were analyzed with both SPM99 and AIBF.

In SPM, the conventional approach was taken: smoothing of data, specification of a temporal design matrix, estimation of parameters and estimation of corrected p-values based on a t-map $SPM\{t\}$. In all experiments, this process was also repeated using *unsmoothed* data, i.e. data to which no spatial smoothing had been applied prior to the analysis. This highlights some characteristics of AIBF in relation to these two types of voxel-based analyses.

The procedure for AIBF was as follows: after specification of the spatial model (Eq. 4.6) and estimation of its parameters for each image (Eq. 4.7), the parameter matrix was

projected back to voxel-space (Eq. 4.19 and 4.21). These images were then analyzed with SPM99 in the temporal domain. The latter step had mainly two advantages. The first was that an existing implementation of a voxel-based approach could be used to compute statistical results for projected AIBF images. The second was that the use of the same temporal analysis and inference for both methods facilitated the cross-validation of AIBF and SPM.

Furthermore, to save computer resources, only a volume of interest (VOI) around the known simulated source was analyzed. This VOI was chosen as a rectangular patch on the flat map encompassing the simulated signal. After projecting this surface patch to voxel-space, a cuboid containing the patch was defined. Search volumes for the estimation of corrected p-values (section 2.2.5) were based on this VOI for all analyses. To facilitate the comparison of AIBF with SPM results, the analysis search volumes for both methods were rendered as similar as possible. Note that one cannot use the same search volume for AIBF (in A -space, Eq. 4.21) and a SPM analysis on smoothed data, because the AIBF search volume in A -space is constrained by definition to the (sparse) AIBF support volume. Smoothed activations might (partially) fall outside this search volume leaving SPM on smoothed data with a disadvantage. Therefore, two different search volumes were used. The first was the AIBF support in A -space and the second was a smoothed version of this search volume, where the kernel was the same as used for smoothing the functional images. This convolution increased the second search volume in comparison to the first search volume by a factor of 3 - 6 (depending on kernel widths). The first search volume was used for the AIBF-analyses in A -space and for the SPM analyses on unsmoothed data. The second search volume was used for the SPM-analyses on smoothed data. The use of two search volumes had the overall advantage that any simulated activation can be detected by SPM on smoothed data making the results of AIBF and SPM more comparable. The disadvantage is that the corrected p-values computed by SPM on smoothed data are slightly increased because of the larger second search volume.

5.2.1 Sensitivity

Introduction

One feature of AIBF is its sensitivity to extended sources confined to the grey matter surface. If the grey matter surface has been correctly identified by the surface reconstruction step and the activation is located close to this surface, then AIBF provides an accurate and sparse spatial model to characterize the observed signal in voxel-space.

In the following experiment, it is shown that AIBF for this simulated data is indeed more sensitive to an underlying non-focal source of activity than a conventional voxel-based method (e.g. SPM99).

Simulations

The simulated data were generated as described in section 5.1.1. The source of activation was defined as a circle with a diameter of 3 mm on the flat map located on the anterior bank of the central sulcus at the presumed hand level. The temporal activation pattern was specified according to a blocked on-off design convolved with a linear combination of two gamma functions to emulate the effects of the haemodynamic response function. The maximum signal intensity of the temporal activation pattern was varied from 0 to 10 % of the global intracortical mean intensity of the functional scans in 1%-steps, giving 10 data sets of 91 images each.

Both AIBF and SPM99 were used to analyze the simulated data sets.

With AIBF, Gaussian basis functions with a distance of 2 mm and FWHM of 2 mm on the flat map were used to construct the model matrix A . The regularization parameter λ was assessed according to Eq. 4.8.

Results

Fig. 5.1 shows the results for each simulated data set. On the left, the maximum t-values for each statistical map are plotted for increasing signal intensities. On the right, the maximum corrected p-values for both methods are shown. At a signal intensity of 2% of the global intracortical mean intensity, AIBF identifies a significant ($p < 0.05$) activation. With SPM99, the signal intensity must be higher than 5 % to result in a significant p-value at the 0.05 level.

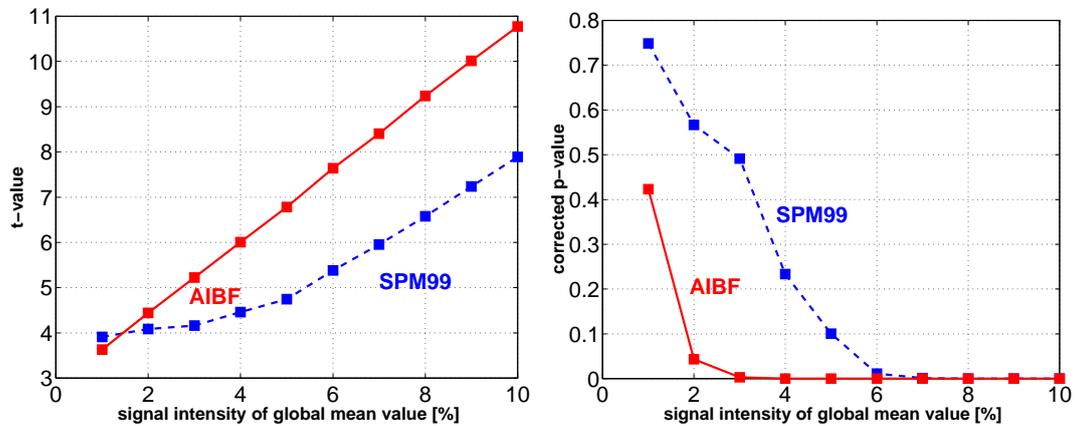


Figure 5.1: Left: t-values computed for 10 data sets with simulated signal using AIBF (red line) and SPM99 (blue line), Right: Same data sets, maximum corrected p-values estimated using AIBF (red line) and SPM99 (blue line)

Discussion

It is obvious from Fig. 5.1 that AIBF outperforms a voxel-based analysis like SPM99 on simulated data sets in terms of sensitivity. The key difference between the two types of analysis is that different forms of feature extraction have been applied to the images prior to the temporal analysis. In a conventional voxel-based analysis like SPM, one typically smooths the images with a spatially invariant three-dimensional convolution kernel. In AIBF, the fitting of basis functions essentially extracts a spatial component from the images, which is confined to the grey matter sheet and is smooth only within this sheet.

Given the assumption that all activity comes from the grey matter sheet or a spatially proximate location, the overall aim is to detect specific spatiotemporal patterns on this surface. The best way to detect these patterns is to find a proper anatomical model describing known features of these patterns and constrain their possible location to the grey matter sheet (c.f. the matched filter theorem).

The conventional way of detecting and localizing activation is to convolve the images with a spatially invariant filter prior to the temporal analysis. The effects of such a smoothing are mainly increases in signal-to-noise ratio and sensitivity, because the interesting signal width is usually greater than the voxel-size, whereas noise also resides in the high spatial frequencies. Note that this smoothing kernel implements its own version of an anatomical model, which is that the observed functional voxel intensities are independent of the underlying anatomy and that activation can be expressed uniformly at all locations in the brain. This is clearly not an ideal model, because it ignores the convoluted nature of the human cortex. If the cortical sheet was a flat surface, convolution with a stationary two-dimensional smoothing kernel prior to the analysis would be a good thing to do, because the kernel exactly reflects the neighbourhood relationships of the underlying neural tissue at any position. However, given the convoluted grey matter sheet embedded in three-dimensional space, a spatially invariant three-dimensional low-pass kernel is no longer informed about the topology of the signal source and cannot be the optimal spatial model.

AIBF, on the other hand, works without any exogenous isotropic smoothing filter. Instead, spatially smooth basis functions located in the grey matter sheet are defined and fitted to each functional observation. Any smooth and extended signal in the grey matter sheet will be captured by this fitting procedure. There are three important features of such a fit: (i) The basis functions follow the grey matter sheet, (ii) the basis functions are informed about the distance measure on the grey matter sheet and (iii) signal from neighbouring structures, like white matter or cerebrospinal fluid (CSF), are attenuated or even removed, because the basis functions have their support largely in grey matter.

Given a smooth activation in the grey matter sheet as simulated in the data, AIBF extracts this spatiotemporal signal from the original image, whereas higher frequency components and signals from outside the support area are precluded from the fitted data. Therefore, the model fitting and the subsequent reprojection of the spatiotemporal parameter matrix B into the voxel-space can be conceived as a spatially variant

anatomically informed smoothing operation, where the smoothing is constrained to the cortical sheet.

As described in 2.2.5, the corrected p-value is approximated by probabilities based on the Euler characteristic in resel space. This renders the corrected p-value in a three-dimensional t-map largely a function of the observed t-value, the search volume, the surface of the search volume and smoothness in the three directions.

The surface area of the AIBF search volume is not much larger than that of the smoothed search volume. This is because the support of the basis functions usually not only covers the grey matter, but neighbouring voxels such as the CSF within a sulcus, although basis function intensities are rather low in these locations. Therefore the AIBF search volume does not have a surface area equal to the total inward and outward side area of the grey matter, but only a fraction of it. Clearly, this fraction depends on the voxel-size and the individual CSF distribution.

Finally, the estimated smoothnesses of the SPM and AIBF are slightly different. The smoothness of the t-map of the SPM analysis is governed by the applied smoothing kernel, whereas the smoothness of the AIBF t-map is a mixture between the width of the basis functions and the regularization used. The latter's effect will be discussed in section 5.2.4. When averaged over all voxels, the measured smoothness in the SPM t-maps was [4.5 4.5 6.8] mm^3 FWHM, which approximates the applied Gaussian smoothing filter of [4 4 6] mm^3 FWHM. In the AIBF t-map, the estimated smoothness was [5.8 4.9 7.6] mm^3 FWHM, i.e. the resel count is smaller in the AIBF than in the SPM t-map further decreasing the corrected p-value in favour of AIBF.

5.2.2 Localization

Introduction

Another feature of AIBF is that the model intrinsically introduces spatial superresolution into the statistical result. This superresolution is because the spatial basis functions are well resolved in the direction orthogonal to the grey matter surface.

This high resolution in the orthogonal direction can be used to better differentiate between sources close to each other in Euclidean image space, but rather distant on the grey matter sheet. A simple example of such a configuration are two juxtaposed sources on opposite banks of a sulcus. If one can resolve such a source configuration with fMRI, this would facilitate the interpretation of statistical maps. With a voxel-based analysis, one would expect that the initial smoothing replaces the two activation sites by a single large one, given that the FWHM of the smoothing kernel is large in relation to the distance between the activations.

In this section, the two activation configuration is analyzed with different AIBF sets and SPM. It will be shown that AIBF can indeed differentiate between the two underlying activations.

Simulations

The simulated data were generated as described in section 5.1.1. For activations, two circles each with a diameter of 3 mm were placed oppositely on the anterior and the posterior bank of the central sulcus at the presumed hand level (Fig. 5.2). Note that the anterior activation is at the location where one would typically find an activation due to a finger opposition paradigm.

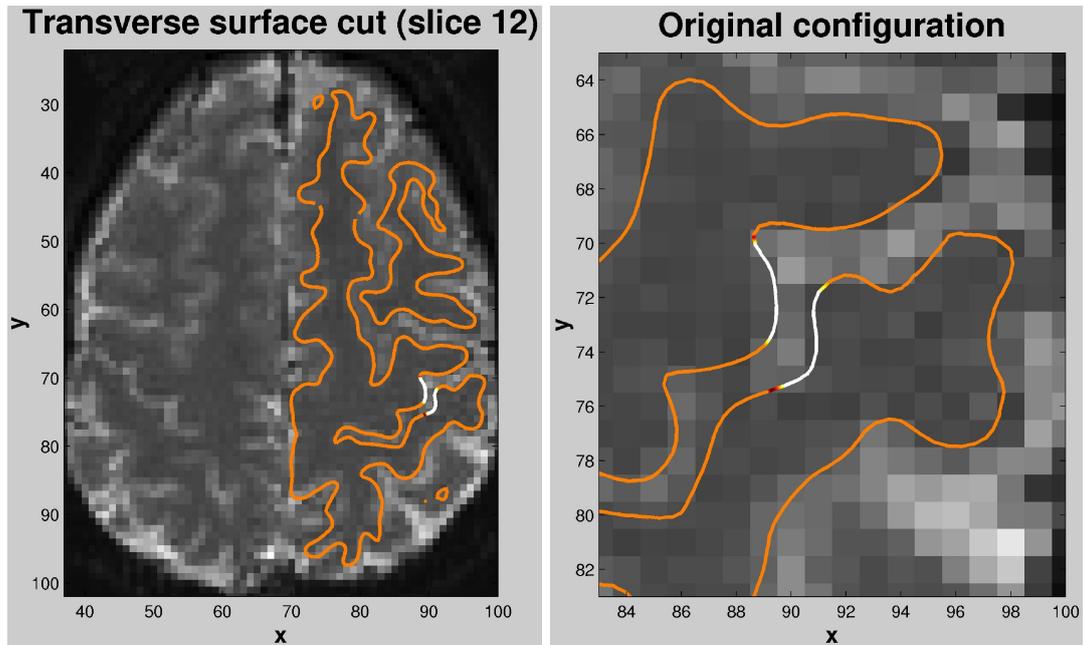


Figure 5.2: Left: Transverse slice and intersecting 3D reconstructed surface at the presumed hand level showing, in white, the location of the simulated activations, Right: (Zoomed-in) region of central sulcus showing the locations of simulated signals

The maximum signal intensity of the temporal activation pattern (see 5.2.1) was 8 % of the global intracortical mean intensity of the functional scans.

With AIBF, two sets of basis functions were used. The first consisted of Sinc basis functions with a distance of 2 mm and width of 3 as defined in Eq. 4.4. The second set comprised Gaussian basis functions with a distance of 2 mm and a width of 2 mm FWHM. The regularization parameter λ was assessed according to Eq. 4.8.

In SPM99, the images were smoothed prior to the analysis with a Gaussian filter ([4 4 6] mm FWHM).

Results

Two sources:

Both AIBF and SPM99 detect the original signal (Fig. 5.3). The best result, in terms

of the smallest p-value, was found by Gaussian AIBF ($1.8e^{-13}$) and with Sinc AIBF a minimum p-value of $1.4e^{-7}$. SPM99 on smoothed data detected the signal with a minimum p-value of $1.5e^{-7}$, on unsmoothed data with a p-value of 0.016. The results in terms of t-maps are shown in Fig. 5.4, Fig. 5.5 and Fig. 5.6.

Both SPM99 on smoothed data and AIBF generate statistical maps, which show evidence of only one cluster, which renders the two different sources undistinguishable and mislocates the maximum of the original signal. SPM99 on non-smooth data seems to detect two separate sources, but only at rather high corrected p-values as compared to the other analyses. The estimated temporal parameter images of both sets of AIBF bear a clearer resemblance to the original signal, i.e. one can differentiate between the two signal sources. The overall pattern generated by the two AIBF sets is similar to the original source configuration (Fig. 5.3).

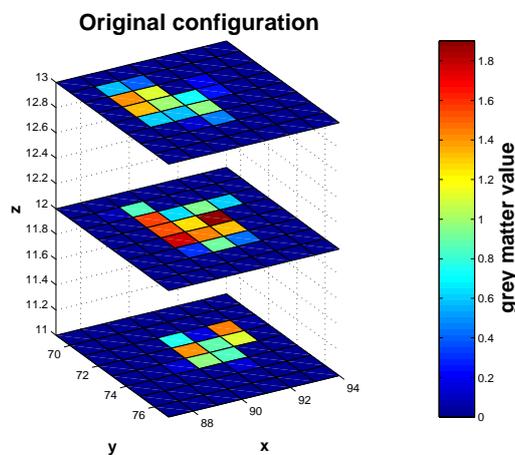


Figure 5.3: Original spatial pattern generated by the two sources configuration in voxel-space

Discussion

In the localization simulations, it was shown that AIBF can differentiate between two underlying activations. This is evident from the temporal parameter images, although not from the t-maps.

Using the statistical maps of SPM on smoothed data or AIBF, one can only identify a single cluster, whose maximum is located half-way between the two underlying sources. In the case of SPM, this result was expected, because the smoothing operation prior to the analysis would merge the two activations into a large one. The AIBF t-maps look similar to the SPM t-map, i.e. one would interpret the result as evidence for one activation between the two banks of the sulcus. The reason for this finding is that the temporal analysis of AIBF is performed on re-projected data such that the temporal noise variance is scaled with the same factor as the underlying parameter. In other words, if one only re-projects only one basis function to voxel-space, one would find a

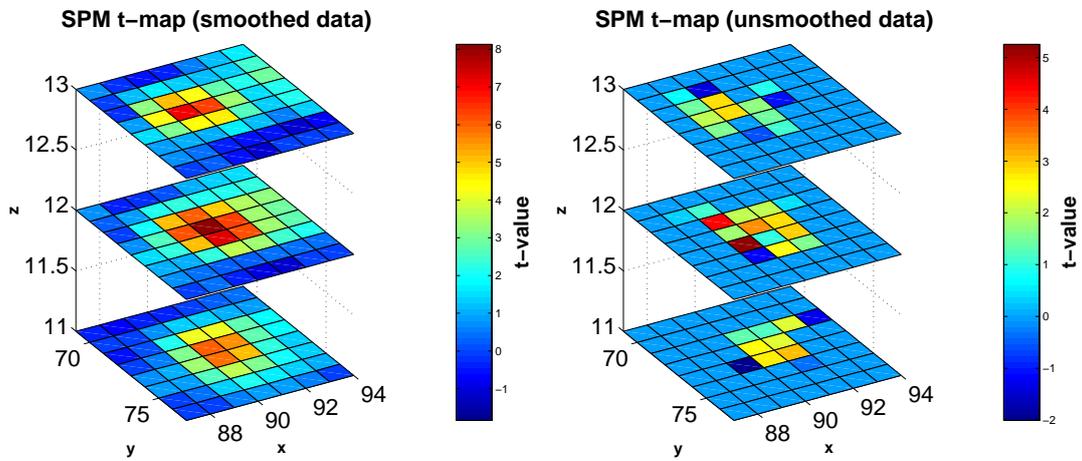


Figure 5.4: Two sources, Left: t-map generated by SPM99 on smoothed data, Right: t-map generated by SPM99 on non-smoothed data

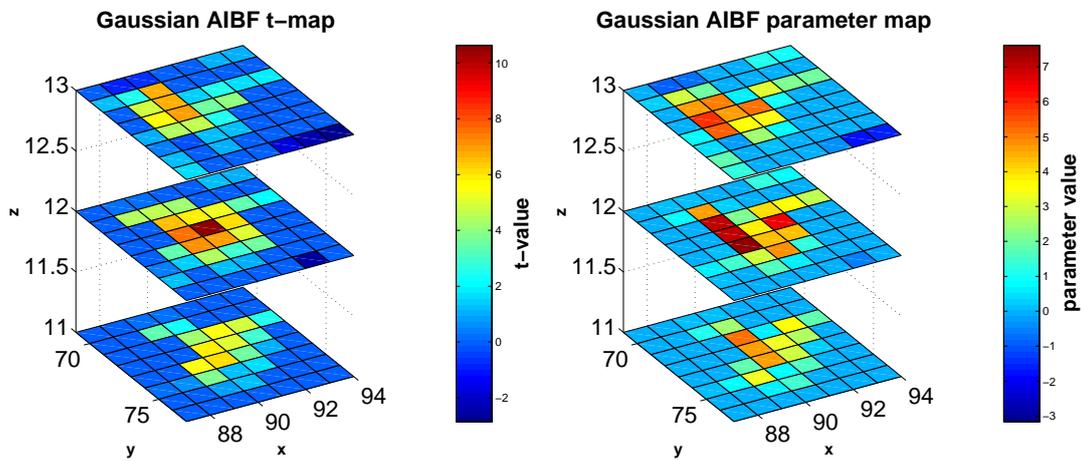


Figure 5.5: Two sources, Left: t-map generated by Gaussian AIBF, Right: parameter map generated by Gaussian AIBF

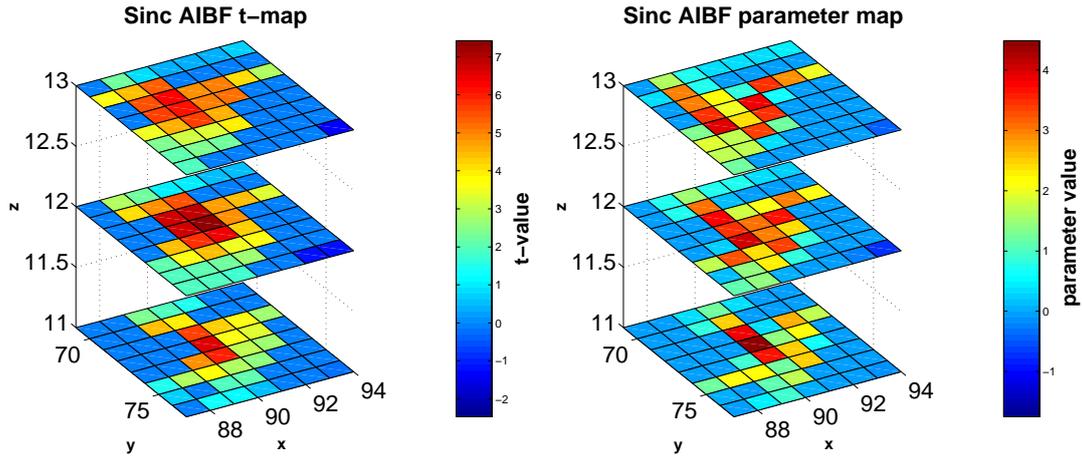


Figure 5.6: Two sources, Left: t-map generated by Sinc AIBF, Right: parameter map generated by Sinc AIBF

constant t-value across the support area of this basis function. If two overlapping basis functions get a high t-value, the maximum t-value will be found in their overlap area, although the estimated temporal parameter is low. This is what can be observed in the AIBF t-maps (Fig. 5.5 (left) and 5.6 (left)).

In contrast, the AIBF temporal parameter maps preserve information about the form of the basis function. Since the basis functions are specified such that the highest basis function values are located within grey matter voxels, the fitted temporal parameters reflect this pattern (Fig. 5.5 (right) and 5.6 (right)).

Although SPM on unsmoothed data seems to be capable of differentiating between the two sources, this comes at the cost of losing sensitivity such that one can differentiate between adjacent activations, but detecting them at all is difficult.

5.2.3 Exploration of parameter space

Introduction

In the two preceding sections, the focus was on the characteristics of AIBF with respect to sensitivity and localization power. The choice of parameters to detect the unknown signal was rather intuitive and not based on any theoretical consideration about the optimal parameter set. In this section, the AIBF parameter space is explored more thoroughly. Given the same underlying activation on the cortical surface, various basis function sets are used to model it. This exploration is presented for all three types of basis sets. The results are characterized by assessing the minimum corrected p-value related to the signal change. Finally, the results are discussed in relation to choosing AIBF parameters for unknown signals.

Simulations

A single configuration was defined in terms of a circle with a diameter of 3 mm, located on the anterior bank of the central sulcus at the presumed hand level. The maximum signal intensity of the temporal activation pattern (s. 5.2.1) was 8 % of the global intracortical mean intensity of the functional scans.

Only AIBF was employed to analyze the data. Three different types of basis functions were used, i.e. (i) Gaussian, (ii) sinc and (iii) discrete cosine sets. For each set, the user-specified parameters were varied to produce a range of different models.

For Gaussian basis functions, the distance between centres ranged between 0.8 and 2.4 mm with a step-size of 0.2 mm, giving nine distances. Each of these distance parameters was applied with a set of basis function widths, ranging between 1.0 to 5.0 mm FWHM with a step-size of 0.4 mm, giving eleven different widths. This resulted in 99 Gaussian basis function sets.

Similarly, for sinc basis functions, the distance between centres ranged between 1.0 and 2.4 mm with a step-size of 0.2 mm, giving eight distances. The widths ranged between 0.1 to 0.8 with a step-size of 0.1, giving eight widths. This resulted in 64 Sinc basis functions sets.

In the case of the discrete cosine sets, the highest spatial frequency was varied, starting from $K_u = 2$ to $K_u = 24$ with a step-size of one and $K_v = K_u$ as defined in Eq. 4.5. All frequencies below the highest frequency were included in a basis set. This resulted in 23 different sets, where the spatial frequencies with respect to the underlying surface S_F ranged between 0.0125 to 0.2674 cycles/mm on the flattened cortical surface.

The regularization parameter λ was determined according to Eq. 4.8.

Results

For each type of basis function (Gaussian, Sinc, discrete cosine), the maximum t-value and corresponding minimum corrected p-value with respect to the simulated signal change were assessed. In the case of Gaussian and Sinc basis functions, this resulted in two-dimensional result maps as a function of the generating parameters distance and width (Fig. 5.7 and 5.8). For the discrete cosine basis functions, the minimum corrected p-value and maximum t-value of each set were plotted as a function of the set order (Fig. 5.9). In all cases, the maximum t-values resulting from the analysis within native parameter space B were computed and plotted.

Discussion

The parameter maps shown in Figs. 5.7, 5.8 and the plot in Fig. 5.9 offer some important insight into the relationship between the user-specified parameters and the width of the underlying activation.

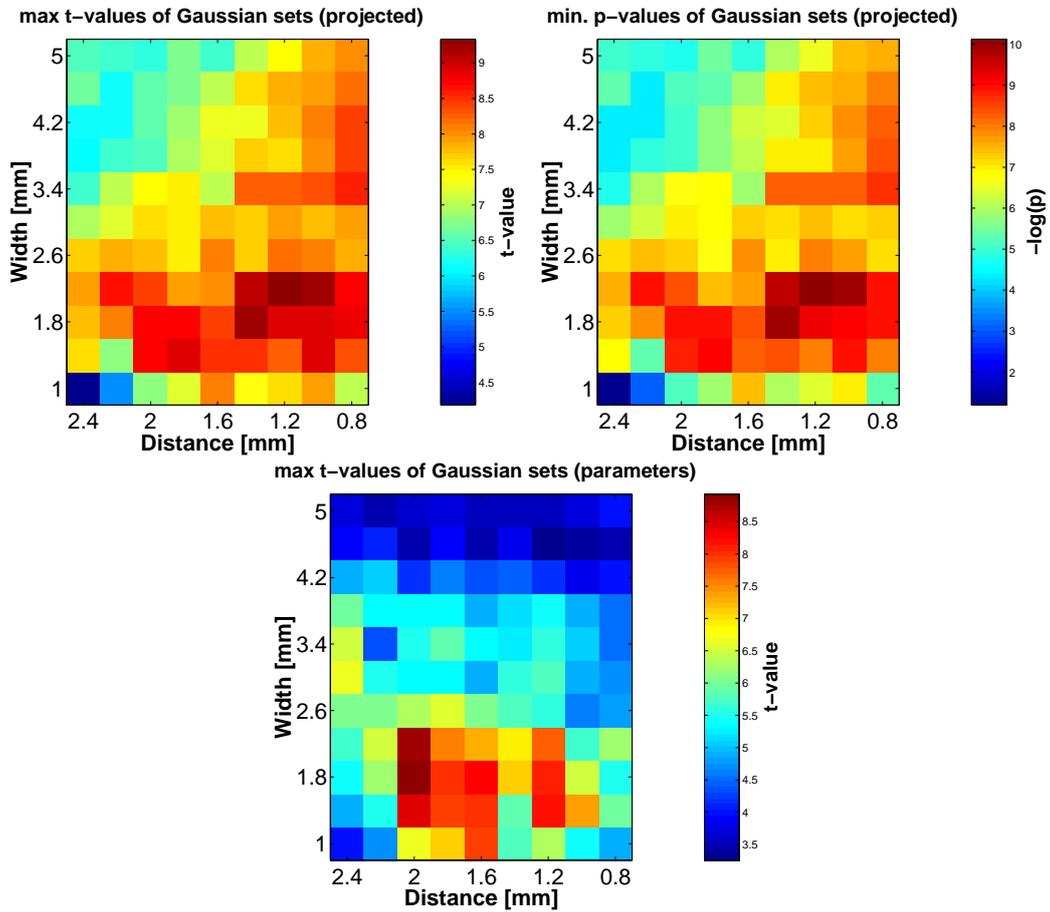


Figure 5.7: Maps of AIBF-results with Gaussian basis functions using a simulated source consisting of a circle with 3 mm diameter. The x-direction encodes the distance between centres and the y-direction the width of basis functions. (Top left) Map of maximum t-values in voxel-space after reprojection (Top right) Map of the negative logarithm of the minimum corrected p-values after reprojection (Bottom) Map of maximum t-values in parameter space

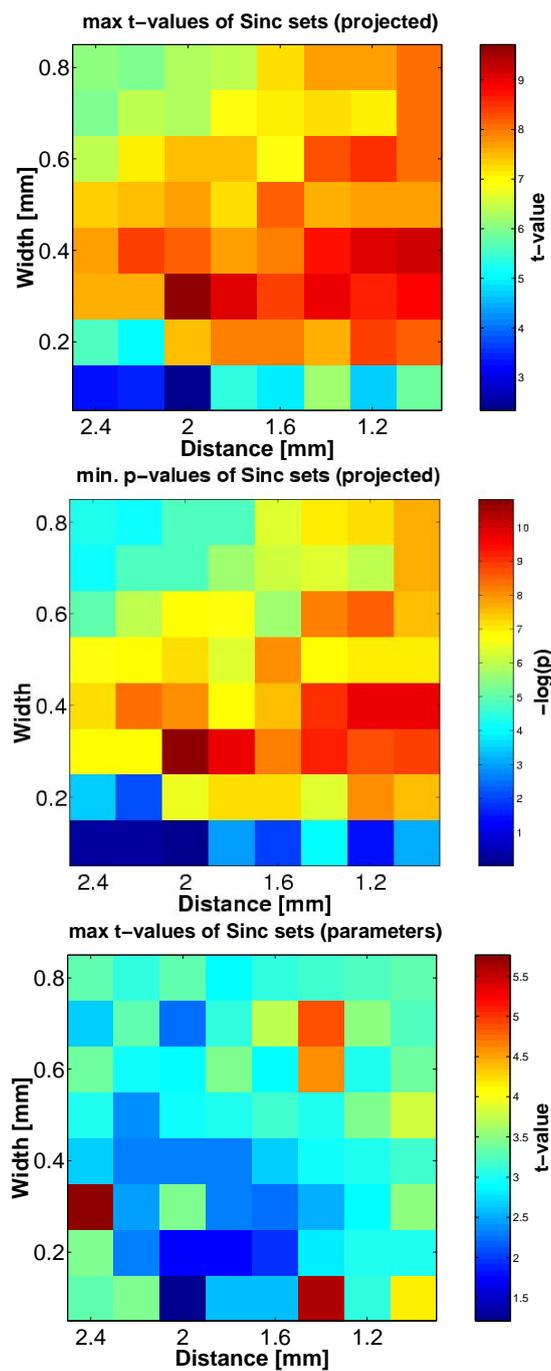


Figure 5.8: Maps of AIBF-results with Sinc basis functions using a simulated source consisting of a circle with 3 mm diameter. The x-direction encodes the distance between centres and the y-direction the width of basis functions. (Top) Map of maximum t-values in voxel-space after reprojection (Middle) Map of the negative logarithm of the minimum corrected p-values after reprojection (Bottom) Map of maximum t-values in parameter space

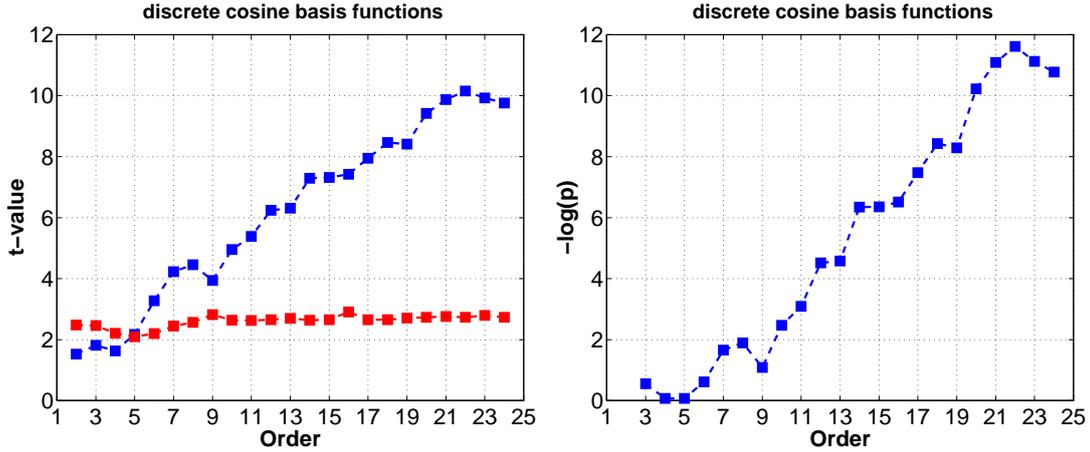


Figure 5.9: Plot of AIBF-results with discrete cosine basis functions using a simulated source consisting of a circle with 3 mm diameter. The x-axis encodes the order of the set. (Left) Maximum t-value in voxel-space after reprojection (blue) and maximum t-value in parameter space (red), (Right) negative logarithm of minimum corrected p-value after reprojection

First of all, it can be seen that the overall maximum t-value assessed for each type of set (Gaussian, Sinc, dct) is between 9 and 10. This corresponds to a minimum p-value smaller than 10^{-10} . These results mean that all three types of sets are suitable for the detection of the underlying focal signal. In comparison to this, SPM computes a maximum t-value of 6.58 on smoothed data and 5.50 on unsmoothed data. This corresponds to minimum corrected p-value of $1.1 \cdot 10^{-4}$ and $6.6 \cdot 10^{-3}$.

It can also be seen from Fig. 5.7 and 5.8 that in the case of Gaussian and Sinc basis functions, the width of the basis functions plays a crucial role with respect to sensitivity. For Gaussian basis functions, the minimum p-values were obtained, when the width, as defined in Eq. 4.2, was between 1.8 and 2.2 mm FWHM. For Sinc basis functions, the optimal width (Eq. 4.3) appears to be 0.3 to 0.4 mm. At this width, good sensitivity can be observed for a wide range of distances. This is a reassuring result, because it means that the exact spacing of the basis function centres does not really matter provided that the width corresponds roughly to the width of the underlying signal. A Gaussian FWHM around 2.0 mm roughly corresponds to the size of the signal. A Sinc function with a width of 0.3 to 0.4 mm is slightly narrower than this Gaussian.

One can also see that small distances appear to lead to better results, even if the width is larger than the signal. This is because the positioning of the basis functions is better with respect to the location of the signal for a denser sampling. Effectively, this means that one should choose not only the optimal width, but also a distance that oversamples the signal.

The choice of parameter gets much more critical if one analyses the estimated parameters B directly in their native space. In the case of Gaussian basis functions, the region of maximal sensitivity is much more confined. The reason for this is that high

sensitivity in parameter space can only be reached if the support of a single basis function covers most of the underlying signal. This is typically the case when the overlap is small and the basis functions are well positioned. In Fig. 5.7 this becomes obvious because the area of high sensitivity in the top right corner in the map of maximum t-values is virtually absent in the map of maximum t-values in parameter space. For Sinc basis functions, the maximum t-values are rather low in comparison to those in the equivalent Gaussian analyses. This might be because Sinc basis functions, with negative side lobes, are a suboptimal basis set to model the simulated signal.

For discrete cosine sets, it can be seen that the t-value and minimum p-value rise until the order of the basis set reaches 22 and then levels off and slightly decreases. The actual maximum side length (Eq. 4.5) of the surface patch was 43 mm, which means for an order 22 the highest frequency basis function has a frequency of 10.5 cycles/43 mm, i.e. 1 cycle/4.10 mm. A half-cycle would then cover 2.05 mm, which is less than the underlying signal width of 3 mm. This lower value is not unexpected, because the spatial form of the underlying signal is not sinusoidal, but rectangular such that higher frequencies are necessary to fit the signal perfectly.

It is reasonable to assume that the overall fit to the signal still improves for higher order sets, but this does not imply better sensitivity. Fitting higher frequencies produce *temporal noise* through overfitting such that the numerator of the ensuing t-value remains the same, but the denominator increases.

The maximum t-values of the discrete cosine sets in parameter space remain very low for all sets, reflecting the fact that the fit of a discrete cosine set is distributed over all basis functions and any single basis function cannot capture the effect.

Clearly, the choice of basis function width or order of a discrete cosine set is a critical parameter. It depends on the prior expectation about the width of the unknown signal. This issue is the same as in conventional voxel-based analyses, where an exogenous smoothing kernel width has to be specified by the user. Usually, in fMRI, the width of the 3-dimensional smoothing kernel will range somewhere between zero and four times the voxel size, depending on the question asked.

For local AIBF (Gaussian and Sinc), one should capture the signal with basis functions that are larger in extent than one voxel with a width equal to the expected signal width. In this thesis, it is assumed that focal signals induce a BOLD effect that has a width of roughly 3 - 5 mm on the cortical surface. The results of this section suggest that one should ensure that the distance between centres oversamples the expected signal width. Given a width of 3 mm, as used for the simulated signal, a Gaussian basis functions set with distance 1.2 mm and width 1.8 to 2.2 mm FWHM would be an appropriate choice.

For discrete cosine sets one should specify an order where the highest frequency is around 1/3 to 1/5 cycle/mm.

5.2.4 Regularization

Introduction

Regularization is a device used in inverse filtering, image restoration and virtually any inverse problem. Prior knowledge or assumptions about the underlying solution are exploited to compromise between fitting the data and the model. The latter incorporates assumptions about the solution. Given that there are many more parameters to be estimated than observations, regularization is one way to resolve the inverse problem, because one effectively restricts the solutions to be associated with a likely model. In this way, regularization can be understood as incorporating soft constraints into the solution. However, in this thesis, regularization is not the key principle. The aim is rather to effectively reduce the number of parameters before fitting the model, i.e. the approach is to project the observations to a low-dimensional subspace using constraints on the anatomy and about the smoothness of the solution. In contradistinction to regularization, this scheme enforces a hard constraint, because a whole subspace of possible solutions have been discounted. The question might be raised, if additional regularization is useful at all, given that the basis function approach already implements some smoothness constraints.

In the following, it is shown that regularization in the context of AIBF can further improve sensitivity.

Simulations

The simulated data was generated as described in sections 5.1.1. One source configuration was specified, which was a circle with a diameter of 3 mm located on the anterior bank of the central sulcus at the presumed hand level. The maximum signal intensity of the temporal activation pattern (see 5.2.1) was 8 % of the global intracortical mean.

Only AIBF models were employed to analyze the data in the spatial domain. All three different types of bases, i.e. (i) Gaussian, (ii) sinc and (iii) discrete cosine were used. The basis parameters were chosen on the basis of the preceding sections as optimal or close to optimum. For Gaussian sets, 1.4 mm separation and 1.8 mm FWHM were used, for Sinc sets 1.4 mm separation and 0.3 mm width and for the discrete cosine set the order was 22.

For each basis set, two techniques to compute λ were applied to two different regularization constraints (section 4.2.6). In the following, application of Eq. 4.8 will be denoted by *estimator 1* and the ReML algorithm by *estimator 2*. The two regularization constraints will be referred to as *constraint 1* (zeroth order regularization in parameter space) and *constraint 2* (first order regularization in voxel-space).

Additionally, minimum p-values (for all basis sets, constraints and techniques) were computed as a function of λ to examine the effects of different degrees of regularization.

Results

Fig. 5.10 plots the minimum p-values for different basis sets and two constraints as a function of λ . The values of λ assessed by the two estimators described in section 4.2.6 are also provided.

Discussion

In this section the effects of regularization on the sensitivity of AIBF was explored. In the following, the relationship among the form of the basis functions, regularization and sensitivity is discussed.

The overlap between adjacent local basis functions is controlled by the width and separation. Decreasing the separation while fixing the width, increases the variance of the estimated AIBF parameters. This is because the model has more *freedom* for fitting the data. One could argue that this overfitting is not really an issue, because one is not concerned with the variance of the estimated parameters, but the estimates themselves. However, an increased variance means that more high frequency components enter the fitted data. If these high frequency components are not signal-related, the sensitivity of the analysis will decrease, because the denominator of the t-value will increase.

This has important implications for the choice of local basis functions and the regularization applied. As concluded, one should choose the separation of basis functions such that the signal is oversampled in relation to its width. As discussed above, the disadvantage is that decreasing separation increases the overlap and sensitivity decreases. This reflects a *trade-off* between the sensitivity to detect changes and the anatomical precision at which these changes are detected. This trade-off can be moderated by a regularization which dampens the parameter estimates and decrease its variance.

As expected, regularization improves the solution with respect to sensitivity, if the width of the basis functions matches roughly the width of the signal and the separation is small. Such a parameter configuration is given by 1.8 mm FWHM and 1.4 mm separation. In (Fig. 5.10), the plot for the Gaussian set shows that without the regularization the minimum p-value is ca. 10^{-4} and experiences a decrease to ca. 10^{-10} , when $\lambda = 1$ and constraint 1 is used. Constraint 2 appears to have deleterious effects in the sense that the sensitivity only increases marginally for local basis functions and decreases for the DCT set.

Note that the smoothness of the solution increases with increasing λ (under both constraints). For example, when $\lambda = 0$, one resel is estimated as 4.04 voxels, whereas for $\lambda = 1$, one resel is 15.49 voxels. If one wants maximum resolution, the obvious choice is to set $\lambda = 0$, i.e. to apply no regularization but at the cost of sensitivity.

If one increases the separation to 2 mm (data not shown), the unregularized solution results in a minimum p-value of $6.27e^{-8}$ (1 resel = 9.61 voxels) and the regularized solution (constraint 1, $\lambda = 1$) improves this to $7.80e^{-10}$ (1 resel = 17.91 voxels). This result indicates that zeroth order regularization at $\lambda = 1$ can be used to find a solution which does not interact with the separation of local basis functions with respect to the

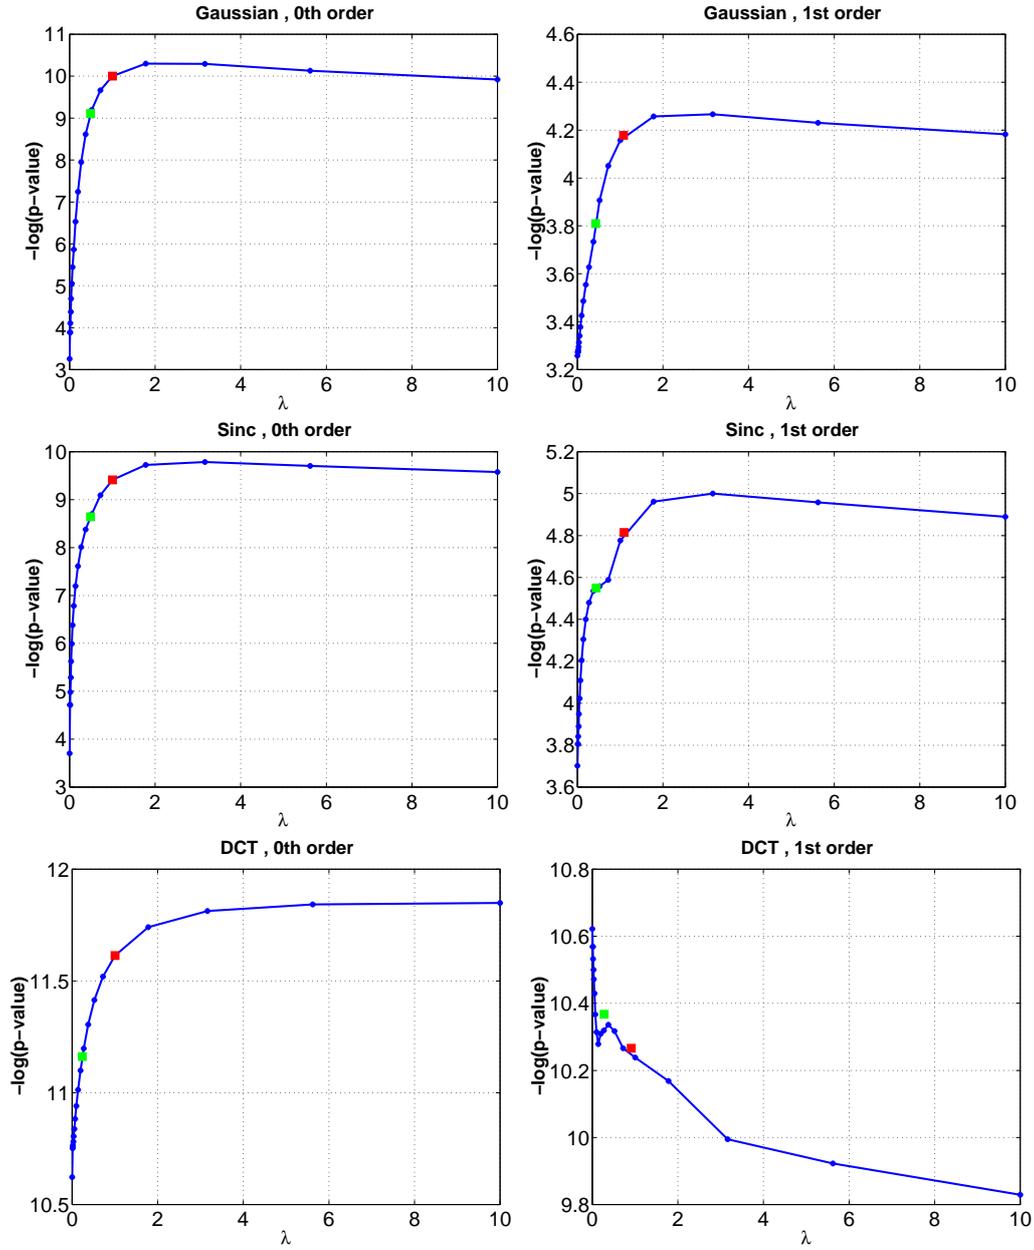


Figure 5.10: Regularization with two constraints and three basis sets. Each graph displays the negative log of the minimum p-values in voxel-space as a function of λ . The red squares mark λ found by using estimator 1, (s. Eq. 4.8). The green squares show the results of estimator 2, the ReML estimator described in section 4.2.6. Top row: Gaussian basis set, middle row: Sinc basis set, bottom row: Discrete cosine set. Left column: constraint 1 (0th-order regularization in parameter space), right column: constraint 2, (1st-order in voxel-space)

sensitivity to the solution.

The results for the Sinc basis functions (Fig. 5.10, middle row) are similar. The same argument about enhancing high frequency components with increasing overlap applies to Sinc functions.

For the global DCT set, the results are slightly different. A DCT set without regularization is more sensitive than local basis function sets. At $\lambda = 0$ the minimum p-value is $2.39e^{-11}$ with 1 resel = 9.3 voxels. At $\lambda = 1$ the minimum p-value is $2.43e^{-12}$ with 1 resel = 12.40 voxels. Regularization with constraint 2 does not improve the sensitivity of a DCT set.

5.2.5 Spatial deconvolution

Introduction

In this section, the application of AIBF to PET data based on simulated activations is presented to illustrate the potential role of AIBF with low resolution data. To a first approximation, the effects of the PET point spread function (PSF) can be characterized by convolution with a Gaussian kernel. In Eq. 4.6, this convolution is implemented by multiplication with a convolution matrix L_I .

In this way, the AIBF model for PET, or any other imaging modality that has a point spread function wider than the voxel-size, is effectively an extension of the AIBF model for fMRI. Furthermore, it will be shown that one can apply an additional convolution to both data and model to render the parameter estimation more robust with respect to registration errors between the grey matter surface and functional data.

Fitting such a convolved model allows back-projection of the estimated parameters not only into the smooth A_L -space (Eq. 4.19), but also into the A -space (Eq. 4.21), i.e. the effects of the PSF and any additional convolution are removed from the fitted data in A_L -space. Effectively, this is a spatial deconvolution of the functional data that is guided by anatomical information. In the following, application of AIBF to simulated PET activation data and the effect of convolution with an additional stationary lowpass kernel is demonstrated.

Simulations

The simulated data were generated as described in section 5.1.2. The PET data used for generating simulated activation data were acquired from a single-subject experiment consisting of twelve scans. A description of this data is in section 6.2.1. The simulated activation varied in an on-off fashion over scans.

A circular source of 16 mm diameter was defined on the cortical surface in the parietal lobe. The maximum signal intensity of the temporal activation pattern (see 5.2.1) was 3 % of the mean grey matter intensity of the functional scans (a small activation). For all AIBF analyses, Gaussian AIBFs with 12 mm separation and 12 mm FWHM were

Type of analysis	correct surface	misregistered surface
AIBF, A -space	0.22	0.68
smoothed AIBF, A -space	0.20	0.16
AIBF, A_L -space	0.44	0.84
smoothed AIBF, A_L -space	0.35	0.29
SPM (unsmoothed data)	0.79	0.99
SPM (smoothed data)	0.53	0.57

Table 5.1: Minimum corrected p-values for peak t-values. Correct surface: There was no registration error between reconstructed surface and functional data. Misregistered surface: The correct surface was shifted by 8 mm in the x-direction. The AIBF parameters were 8 mm separation and 8 mm FWHM. The *smoothed* AIBF set involved applying an isotropic Gaussian smoothing to the data and the model. The SPM analyses were performed on smoothed (12 mm FWHM in each direction) and on unsmoothed data.

used. These data were analyzed with two different convolution matrices. The first was a convolution matrix L_I ([8 8 8] mm FWHM) to emulate the effects of the PET PSF. In the second simulation, both the data and model were further multiplied with a convolution matrix L_E of [12 12 12] mm FWHM. In the following, the first simulation is referred to as the AIBF model, and the second is the smoothed AIBF model.

These two simulations were repeated, where surface misregistration was emulated by shifting the reconstructed surface by 8 mm (4 voxels) in the positive x-direction. In a conventional SPM-analysis, a Gaussian smoothing filter of [12 12 12] mm FWHM was employed.

Results

In table 5.1, the minimum p-values of all analyses are listed. In Fig. 5.11, the simulated signal, the structural MRI and the p-values for each analysis are displayed through a coronal slice.

Discussion

In this section, it has been shown how AIBF can be applied to PET data by using an extension of the proposed AIBF model for fMRI. This obtains by multiplying the model with a convolution matrix (Eq. 4.6), which emulates the effect of the PET point spread function. Furthermore, it has been demonstrated that additional isotropic smoothing can be applied to both the data and model to render the AIBF model more robust with respect to registration errors between the anatomical and functional data.

In tab. 5.1, the minimum p-values for peak t-values are listed for various analyses. One can see that the AIBF analyses based on the correct surface are more sensitive than a SPM analysis. Interestingly, the sensitivity is roughly the same in both AIBF analyses.

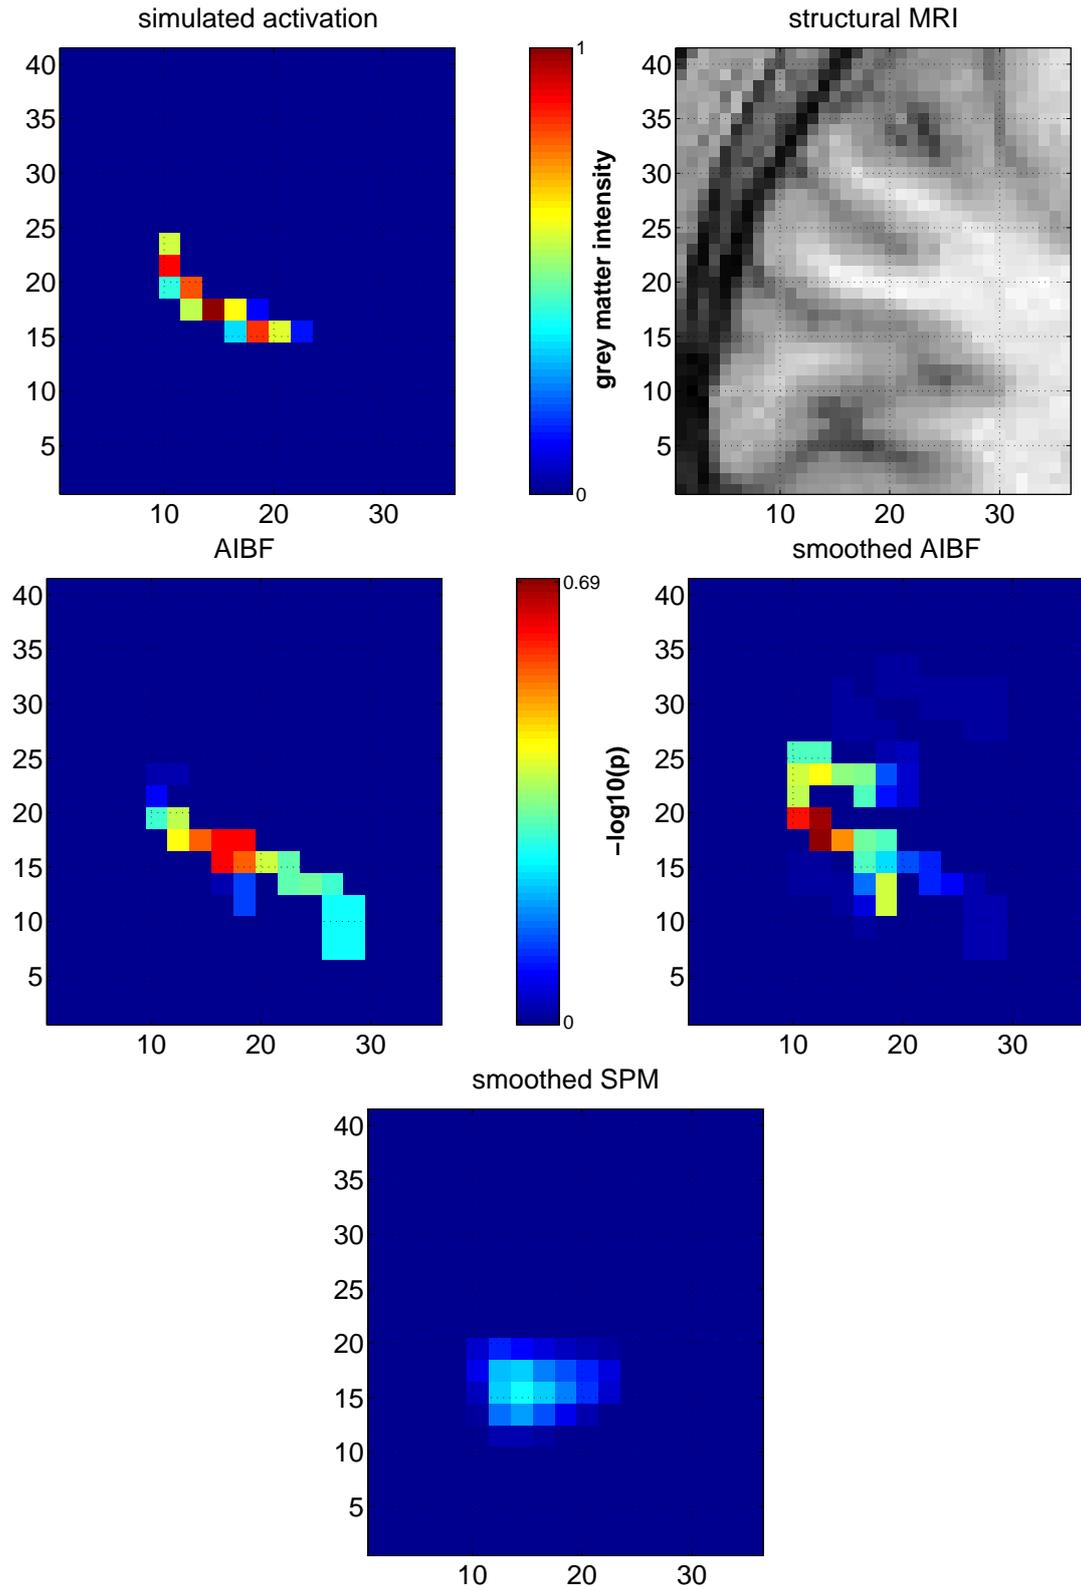


Figure 5.11: Maps showing the negative log of p-values for three analyses. Top row: simulated activation and structural MRI, middle row: AIBF analyses (parameters described above) in A -space, bottom row: SPM analysis

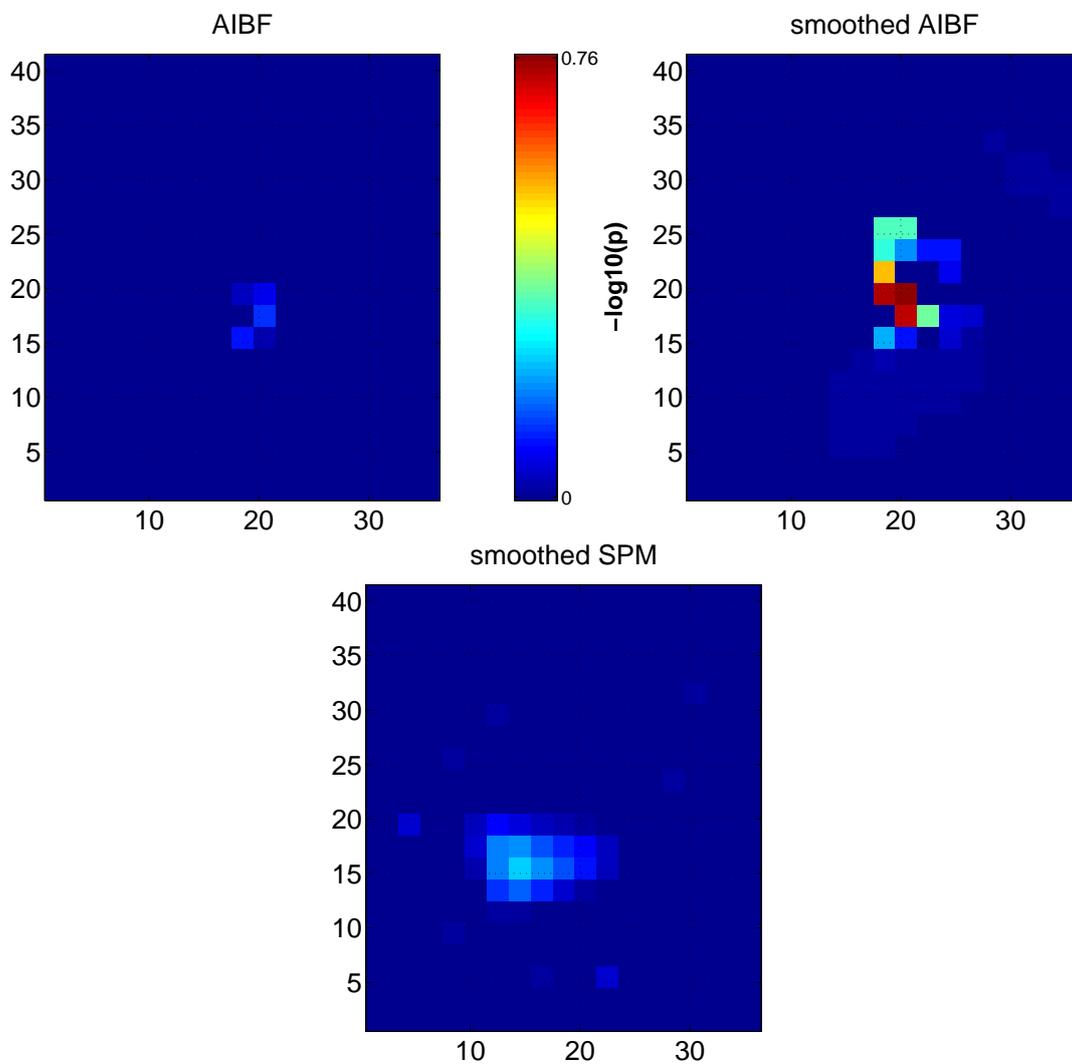


Figure 5.12: Maps showing the negative log of p-values for three analyses, where the AIBF analyses are based on a misregistered surface. Top row: AIBF analyses (parameters described above) in A -space, bottom row: SPM analysis

Although this is a reassuring result, it should not be expected for all data, because sensitivity depends critically on the width of the underlying signal.

When using a misregistered surface, AIBF without additional smoothing is obviously no longer able to capture all of the underlying signal and sensitivity decreases. Sensitivity is restored if additional smoothing is employed. The minimum p-values are even slightly lower than with the correct surface.

Note that the minimum p-value of the SPM analysis on unsmoothed data decreases, when the translated surface is used. This is because the same search volume was used as in the AIBF analysis in A -space, which is essentially a narrow ribbon following the cortical sheet. In the analysis where the translated surface was used, the voxel location showing maximum evidence of an activation was simply excluded from the search volume. Obviously, this means that constraining the search volume has the advantage that the sensitivity to detecting an activation increases within the search volume, but this kind of constraint is not very robust with respect to an ill-informed choice of the search volume.

The minimum p-value of the SPM analyses on smoothed data varies slightly which is because of different estimates of the resel count within the two search volumes.

In Fig. 5.11, one can observe the effect of the additional smoothing filter on the AIBF model fit. In the middle row, two coronal slices of the two AIBF maps (correct surface) are shown. The image on the right hand side shows enhanced evidence for an activation at those locations that are close in image space to the true underlying activation. Although no activation was placed at e.g. position (15, 24), the smoothed AIBF analysis finds a p-value of 0.5 there, because this location is quite close in image space to the true activation.

The applied shift of 8 mm to the surface to simulate a registration error between structural and PET data is rather unlikely to occur for single subject data. From experience and simulated studies (Kiebel et al., 1997), the registration error between PET and MRI image is at most 2 mm. However, there are two important cases where registration errors on this scale play a role.

The first is if the surface reconstruction is partially incorrect. For instance, this could be caused by low contrast in parts of the structural MRI between grey and white matter tissue. The temporal lobes are a region of the brain where this can happen because of its off-center location and the thin white matter tracts in this region that incur partial volume effects. The principled solution is obviously to improve the surface reconstruction, but it is worthwhile to keep in mind that a second option is to smooth both model and data to ensure robustness of the solution. Although this degrades image resolution, it is better to have a suboptimal model as a first approximation than to have a model highly susceptible to slight errors. The same applies to fMRI data, where EPI distortions can cause wrong registration between structural and functional MRI data.

The second reason to apply a smoothing filter to both data and model is given by multi-subject studies. As it will be demonstrated in the next chapter, the basic idea

is to stereotactically normalize the functional data from different subjects and choose a canonical surface, which can be e.g. the normalized surface of one subject. In this thesis, an image-based normalization was used ([Ashburner and Friston, 1999](#)) such that surface registration errors between different subjects are possible. However, as it was demonstrated in this section, convolution with a lowpass filter can be used to overcome the registration error between canonical surface and individual activation. This issue will be further discussed in the next chapter.

Chapter 6

Applications

In the previous chapter, AIBF was validated and tested on simulated data. Some guidelines were derived pertaining to what parameter combinations are appropriate to model data. In this chapter, AIBF is applied to real fMRI and PET data. For fMRI, data from five subjects were analyzed in a series of single-subject studies and in a group study. For PET, AIBF is applied to two group studies. It is shown that AIBF offers a way of controlling the balance between resolution and sensitivity. In particular, two applications are studied in detail. The first is the case where one is interested in enhancing the spatial resolution by employing an accurate anatomical model of the underlying individual neuroanatomy. This is most useful in the analysis of single subject fMRI statistics. The second application employs AIBF for group studies of fMRI or PET data, where one can increase the sensitivity at the expense of resolution tangential to the cortical surface but retain high resolution normal to it.

6.1 Functional magnetic resonance data

In this section, AIBF and SPM are applied to fMRI data. Data from five subjects are used to demonstrate how AIBF can be used for three kinds of analyses. Two models are used for single subject analyses, where the first analysis achieves good resolution on the cortical surface. The second AIBF model puts more weight on sensitivity than resolution. The third model is used to analyze the five data sets as a group study, where one assumes that the anatomical variance can be decomposed into two error components, one is the residual transverse error along the position of the canonical surface relative to the true surface and the second the translational misregistration along or normal to the surface.

The paradigm is fingertapping, a common experimental *test* in fMRI, because it evokes robust activations in the contralateral primary sensorimotor cortex. It is not clear whether the activation detected with fMRI in finger opposition tasks is truly a picture of what happens at the neuronal level. However, it is known that finger opposition tasks generate strong evidence for a BOLD effect at the so called *hand knob*, which is a macroscopic anatomical feature easy to identify on high-resolution structural MR

images (Yousry et al., 1997). This hand knob is located in the anterior bank of the central sulcus at $z \sim 48$ in Talairach coordinates. Given this robust paradigm, the primary interest in this study was to compare different models with respect to sensitivity and resolution for activations located around the hand knob.

6.1.1 Single subject studies

Data acquisition and design

Data from five right-handed subjects were analyzed, where each subject performed a finger opposition task. Subjects opposed the middle finger and thumb of the right hand at a self-paced frequency of 0.5 Hz. All subjects were right-handed.

The functional data (EPI, gradient echo) were acquired on a clinical 1.5 T Siemens Vision unit (Siemens GmbH, Erlangen) with TR = 168 ms, TE = 79 ms, flip angle = 90° . Each image consisted of 128×128 voxels, 20 slices, transverse orientation, voxel size $1.8 \times 1.8 \times 3$ mm³. For each subject, the position and orientation of the slices were chosen to cover the brain volume superior to the corpus callosum, in particular the primary sensorimotor hand area.

The design was blocked and each epoch consisted of seven scans. Rest and activation epochs were alternated. 94 scans were acquired, where the first three scans were discarded to allow for a steady state of the transverse magnetization to be obtained. The first condition was rest, giving seven rest and six activation epochs.

Analysis

The focus of interest, in this study, was activation in the left central sulcus at the presumed hand level. Therefore a rectangular VoI on the reconstructed flattened surface was chosen to cover roughly 3600 mm² of the flattened grey matter surface in the left hemisphere encompassing the central sulcus. The exact VOI varied slightly over subjects. All data were normalized to facilitate a later group study and comparison of the single subject results within a common anatomical frame.

SPM99 was used on non-smoothed and smoothed functional data, where the Gaussian smoothing kernel had a FWHM of [4 4 6] mm, i.e. roughly twice the voxel-size in each direction.

The first AIBF analysis aimed at maximizing the effective resolution of the results. Gaussian basis functions were used to generate the spatial model. The parameters were 1.5 mm separation and 1.5 mm FWHM in each direction on the cortical surface.

The goal of the second AIBF analysis was to enhance the sensitivity of the AIBF model. To do this, the data were smoothed prior to spatial modelling with a Gaussian filter of [2 2 3] mm FWHM. The model was smoothed with the same filter kernel. The Gaussian basis function parameters were 3 mm separation and 3 mm FWHM.

Subject	SPM unsmoothed		SPM smooth	
1	$1.1e^{-6}$	-38/-26/51	$1.2e^{-11}$	-40/-22/54
2	$4.7e^{-11}$	-44/-24/63	$1.1e^{-11}$	-42/-24/60
3	$2.2e^{-16}$	-46/-27/54	$2.2e^{-16}$	-44/-27/69
4	$2.7e^{-12}$	-42/-22/54	$2.7e^{-14}$	-44/-22/69
5	$2.8e^{-15}$	-47/-17/60	$2.2e^{-16}$	-49/-24/60

Table 6.1: Minimum corrected p-values for maxima and their locations in Talairach space

Subj.	AIBF(resolution)		AIBF (sensitivity)	
1	$1.5e^{-9}$	-38/-24/51	$2.6e^{-13}$	-38/-24/51
2	$5.1e^{-12}$	-40/-26/63	$3.8e^{-14}$	-37/-24/60
3	$2.2e^{-16}$	-46/-29/54	$2.2e^{-16}$	-44/-24/54
4	$5.6e^{-13}$	-40/-27/51	$2.2e^{-14}$	-40/-24/66
5	$2.2e^{-16}$	-49/-26/63	$4.4e^{-16}$	-47/-20/63

Table 6.2: Minimum corrected p-values for maxima and their locations in Talairach space

The search volumes of AIBF and SPM on unsmoothed data were given by the volume of support of the basis functions in voxel-space. The search volume of SPM on smoothed data was a smoothed version of the AIBF search volume, where the smoothing filter applied was the same as used prior to the SPM analysis, i.e. [4 4 6] mm FWHM. Note that the estimator of the corrected p-values is a function of search volume (section 2.2.5).

Results

In all five subjects and across all analyses, strong evidence for an activation was found in the primary sensorimotor area in the left hemisphere. In tables 6.1 and 6.2, the minimum corrected p-value for each subject is listed for each method. Note that the minimal corrected p-value for SPM99 is $2.22e^{-16}$ because of a ceiling effect in its computation.

To further characterize the distribution of activation along the z-axis, the maximum t-value for each plane in normalized space was plotted against the z-height (Fig. 6.1).

In Fig. 6.2, the t-maps of subject 4 around height z=51 mm are shown for each analysis. Subject 4 showed a *typical* activation, representative of the other subjects.

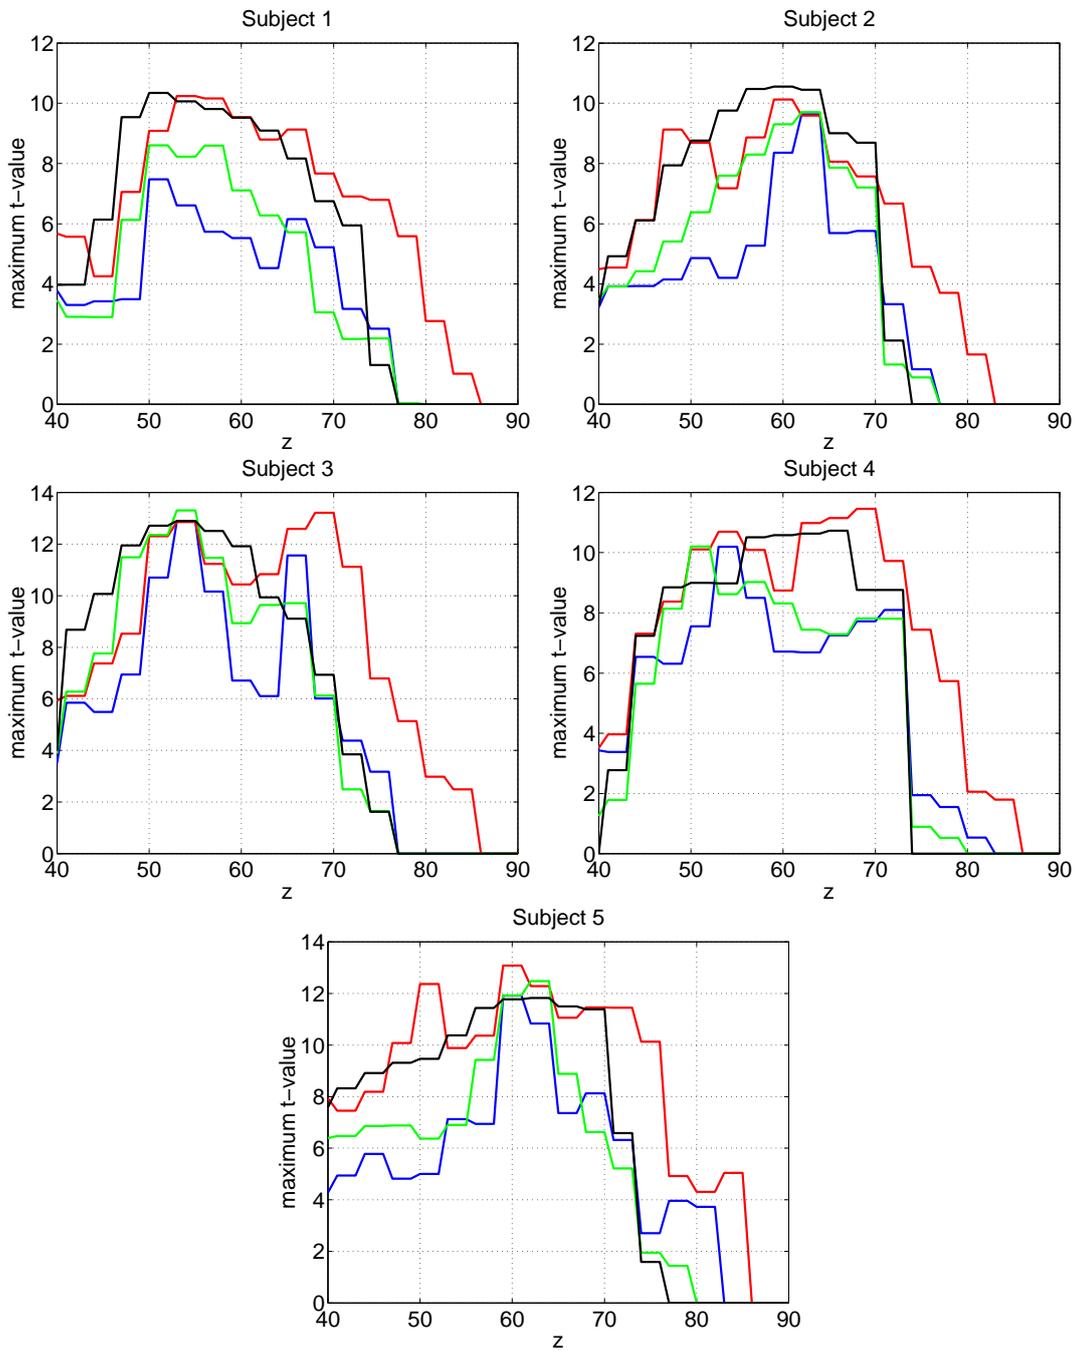


Figure 6.1: t-values as a function of normalized z-height for five subjects. (Blue) Maxima of SPM on unsmoothed data, (Red) Maxima of SPM on smoothed data, (Green) AIBF (resolution), (Black) AIBF sensitivity

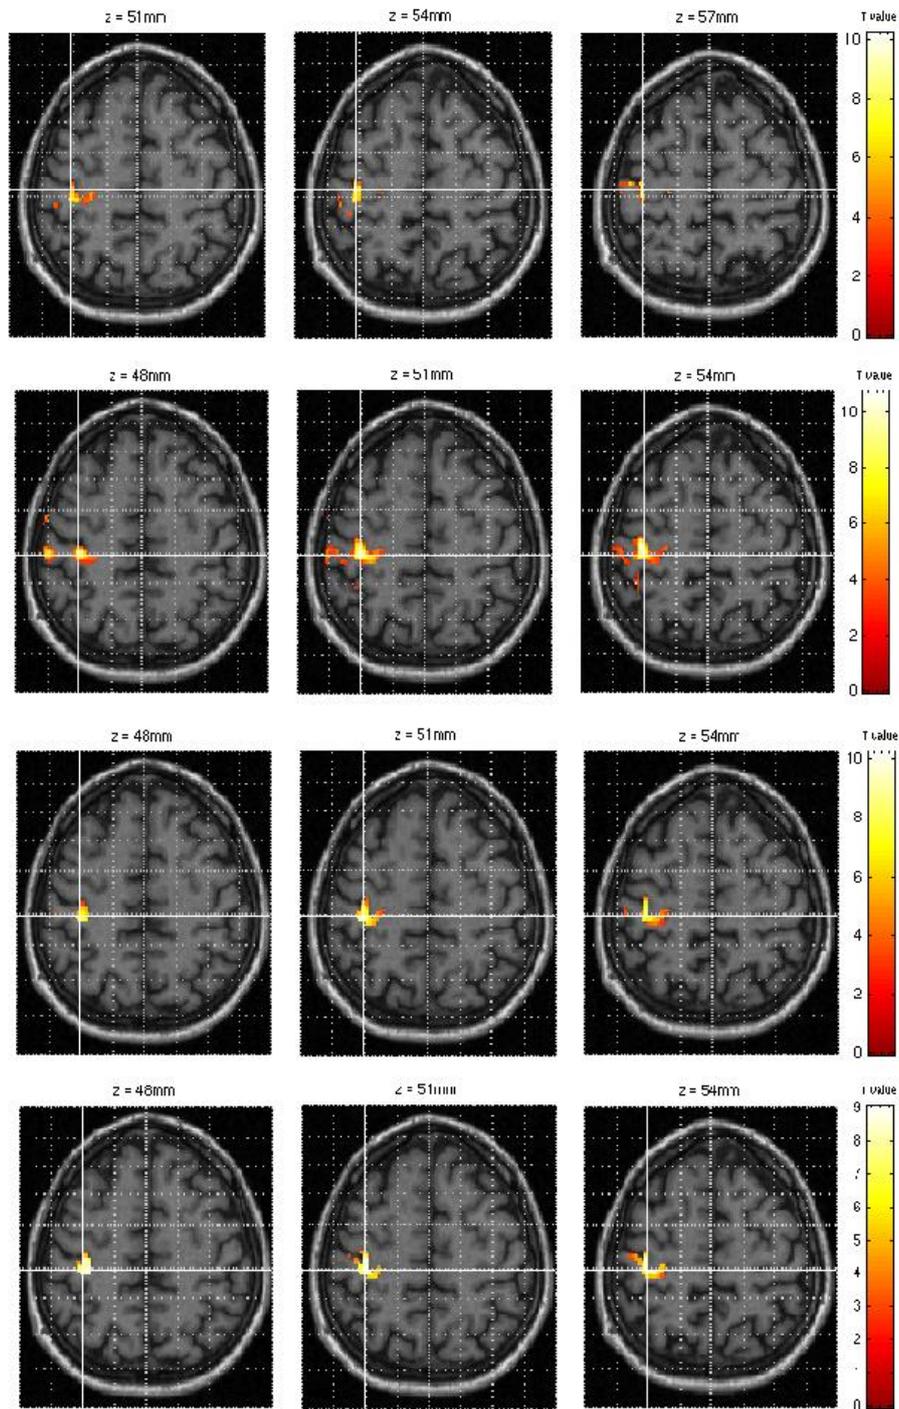


Figure 6.2: t-maps at $z=51$ mm for 4 different analyses on subject 4, thresholded at $p(t) < 0.001$ (uncorrected). First row: SPM on unsmoothed data, second row: SPM on smoothed data, third row: AIBF (resolution) and bottom row: AIBF (sensitivity)

Discussion

In this section, SPM and AIBF were used to analyze five single subject data sets. Gaussian AIBFs were applied with two different parameters, giving two spatial models.

The first AIBF model captures low frequency components located in the grey matter sheet such that signal components from outside the grey matter sheet are suppressed. This results in a statistical map with evidence for signal changes that have two spatial characteristics: (i) they are confined to the grey matter sheet and (ii) they consist of spatial low-frequency components distributed along the direction of the cortical sheet. Note that the first constraint is subject to the resolution of the measured data and the anatomical validity of the model, i.e. through partial volume effects and potential surface misregistration, signal components originating from extra-parenchymal space can influence the estimates. Nevertheless, this particular AIBF model extracts signal components such that sensitivity for smooth BOLD effects located in grey matter is enhanced while other signal sources are attenuated.

The second AIBF model weakens the constraint about the location of signal and explicitly allows sources close to the grey matter to contribute to the solution. This is a consequence of smoothing the data and the AIBF model with an isotropic filter prior to the parameter estimation. The estimated parameters are then projected to the voxel-space defined by the unconvolved model and analyzed over time. The key point here is that the AIBF step can also be understood as a projection of the functional data onto the grey matter surface followed by tangential smoothing along the cortical surface.

The most sensitive results, with respect to an activation within the VOI, were found for SPM on smoothed data and the second AIBF model. The first AIBF model is more sensitive than SPM on unsmoothed data. Note that the p-values are only comparable between the smoothed SPM and the AIBF analyses in a limited way because of the different search volumes. As discussed in section 5.2, the search volume of SPM on smoothed data is roughly five times larger than the search volume of SPM on unsmoothed data or of both AIBF models. The smoothness is lowest for SPM with unsmoothed data, where one resel is ca. 4 voxels. One resel with smoothed data is about 20 voxels. One resel for the first AIBF model is roughly 40 voxels and about 150 voxels for the second AIBF model.

Resolution can be seen from two points of view in spatial modelling. In voxel-based analyses of single subject data, the resolution is governed by the reconstruction filter and the exogenous isotropic smoothing filter. Fitting AIBF models can be thought of as applying a non-stationary smoothing filter where the characterization of the smoothness in terms of a Gaussian three-dimensional kernel is an abstract one. This can be seen in Fig. 6.2, where it is clear that the AIBF t-map is very smooth tangential to the cortical surface, but rough in the orthogonal direction. This means that AIBF, in this example, resolves less acutely than SPM along the cortical surface. But there are two important advantages. The first is that AIBF attenuates or removes confounding signal sources. In other words, one can increase the smoothness (and sensitivity) along the surface without incorporating unwanted signal sources that lie outside the cortex. The second advantage is that the resolution of AIBF might be degraded along the surface, but it is

excellent in image space. For instance, if two activations are close in image space, but remote along the surface, AIBF can resolve them well. Note that this means that one can increase the smoothness of the AIBF statistical maps using wider basis functions without smearing activations together.

The t-maps from all methods show one activation cluster elongated in the z-direction, from roughly 48 mm to 70 mm superior to the AC-PC (anterior commissure-posterior commissure) line (Fig. 6.1). One explanation for this observation might be that a large area of the primary sensorimotor cortex is activated. Another explanation is that an activation, roughly at hand level (say $z = 48$ to 54 mm) accounts for the ventral component and the dorsal part of the activation is caused by a BOLD effect from draining veins. This explanation is supported by the fact that the most dorsal part of the activation is clearly located outside the grey matter. This effect is most evident in the SPM analyses of smoothed data. There the most dorsal part of the activation is clearly located in CSF and reaches $z = 70$ to 85 mm, whereas the AIBF approach locates the most dorsal part roughly 10 mm more ventrally. Note that SPM using unsmoothed data works on masked images, where the mask conforms to the support of the AIBF basis functions. This means that if there are sources outside the grey matter, SPM with unsmoothed data, as used here, would not show it.

Interestingly, if one characterizes the activation by a single maximum (peak-value) one can see from Tab. 6.1 that SPM on smoothed data locates the minimum p-value in all subjects (except for subject 1) somewhere in the dorsal part of the activation. This is especially marked in subjects 3 and 4, where SPM on smoothed data finds a minimum that is located very dorsal to the presumed hand level and outside the grey matter. This is not the case for AIBF and SPM with unsmoothed data. Both methods locate the minimum p-value in subjects 3 and 4 adjacent to the hand knob.

For all other subjects, the locations of the maxima found by the three methods are similar. Visual inspection of the overlay of each statistical map showed that all maxima identified by AIBF were indeed close to or located within grey matter. In subjects 2 and 5, all methods located the minimum p-value around $z = 60$ mm. In Fig. 6.1, it can be observed that SPM with smoothed data finds a local maximum t-value around $z = 50$ mm. These are indeed close to the hand knob, but outside grey matter. The first AIBF analysis discards these activations, whereas the SPM analysis on smoothed data and the second AIBF analysis estimate a less significant p-value around this location.

6.1.2 Multiple subject study

The same five fMRI data sets of the previous section were also analyzed as a group study. In this application, the AIBF model is based on a canonical surface, which is regarded as representative of all the individual surfaces. The important point, made in this section, is that AIBF offers a new way of inter-subject pooling. In conventional analyses, a large stationary smoothing filter is applied to overcome the anatomical variance, i.e. the difference in functional anatomy among subjects. The problem of anatomical variance is that functional localization does not seem to depend on macroscopic structural landmarks except in primary areas, where some relationship between

function and structure has been established (e.g. the hand knob feature). In other words, one cannot predict the exact location of functional areas given only structural information. This fundamental problem cannot be solved by AIBF, but it offers a way to effectively project the sources onto a canonical surface in which their overlap is ensured by smoothing within this surface. This can be implemented by smoothing the data and the AIBF using a Gaussian three-dimensional filter as in the previous section. Here the aim is to render the method robust to intersubject variations in sources (as opposed to misregistration errors). After fitting the data to a smooth AIBF model, one projects the estimated parameters back to the (unsmoothed) voxel-space and analyzes the resulting images over time. The advantage of this approach is that the data is effectively smoothed along the canonical surface with a non-stationary filter, which is a convolution along the surface and a deconvolution transverse to it. The *deconvolution* component can be thought of as *undoing* the smearing due to intersubject variability.

Analysis

Two SPM and two AIBF analyses were performed. For the SPM analyses, the data were smoothed with a Gaussian filter of [6 6 9] mm FWHM and [8 8 12] mm FWHM. These two filter kernels were found to be optimal for this data set among Gaussian kernel widths ranging from 0 to 18 mm FWHM in each direction. The first AIBF parameter set was Gaussian with 3 mm separation and 3 mm FWHM on the surface combined with a Gaussian L_E of [2 2 3] mm FWHM. The second was a Gaussian AIBF set with 3 mm separation and 3 mm FWHM and a Gaussian L_E with [4 4 6] mm FWHM.

Results

In Fig. 6.3, the t-value maxima of each slice for each analysis are plotted against z-height. Fig. 6.4 shows the thresholded t-maps of all four analyses at z-height 63 mm.

Discussion

The most telling result is displayed in Fig. 6.3. One can see that all analyses find highly significant t-values between z-height 55 mm and 68 mm above the AC-PC line. The AIBF analyses tender higher t-values than any of the SPM analyses. The highest t-value (and most significant) result is given by AIBF applied to smoothed data, where the three-dimensional filter kernel implemented by L_E had [4 4 6] mm FWHM. Generally the most sensitive result is obtained if one specifies a filter which has roughly the shape of the elongated activation found in the single subject analyses. This is most efficiently implemented with AIBF, because, with an anatomically informed basis, one can model smooth components following the cortical surface. The three-dimensional Gaussian filter controls the distance from which signal from locations close to the cortical surface enter into the estimate. AIBF is, for these data, ideally suited to extract smooth spatial components subtending the activation.

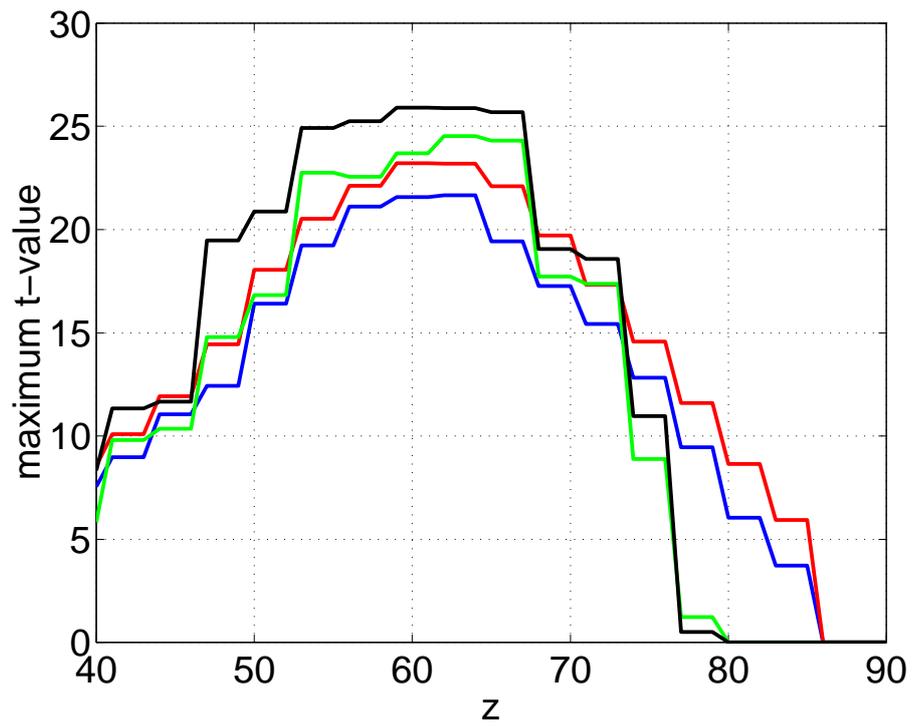


Figure 6.3: t-value maxima as a function of z-height (vertical position) for five subjects. (Blue) Maxima of SPM, [6 6 9] mm FWHM, (Red) Maxima of SPM [8 8 12] mm FWHM, (Green) AIBF [2 2 3] FWHM, (Black) AIBF [4 4 6] FWHM

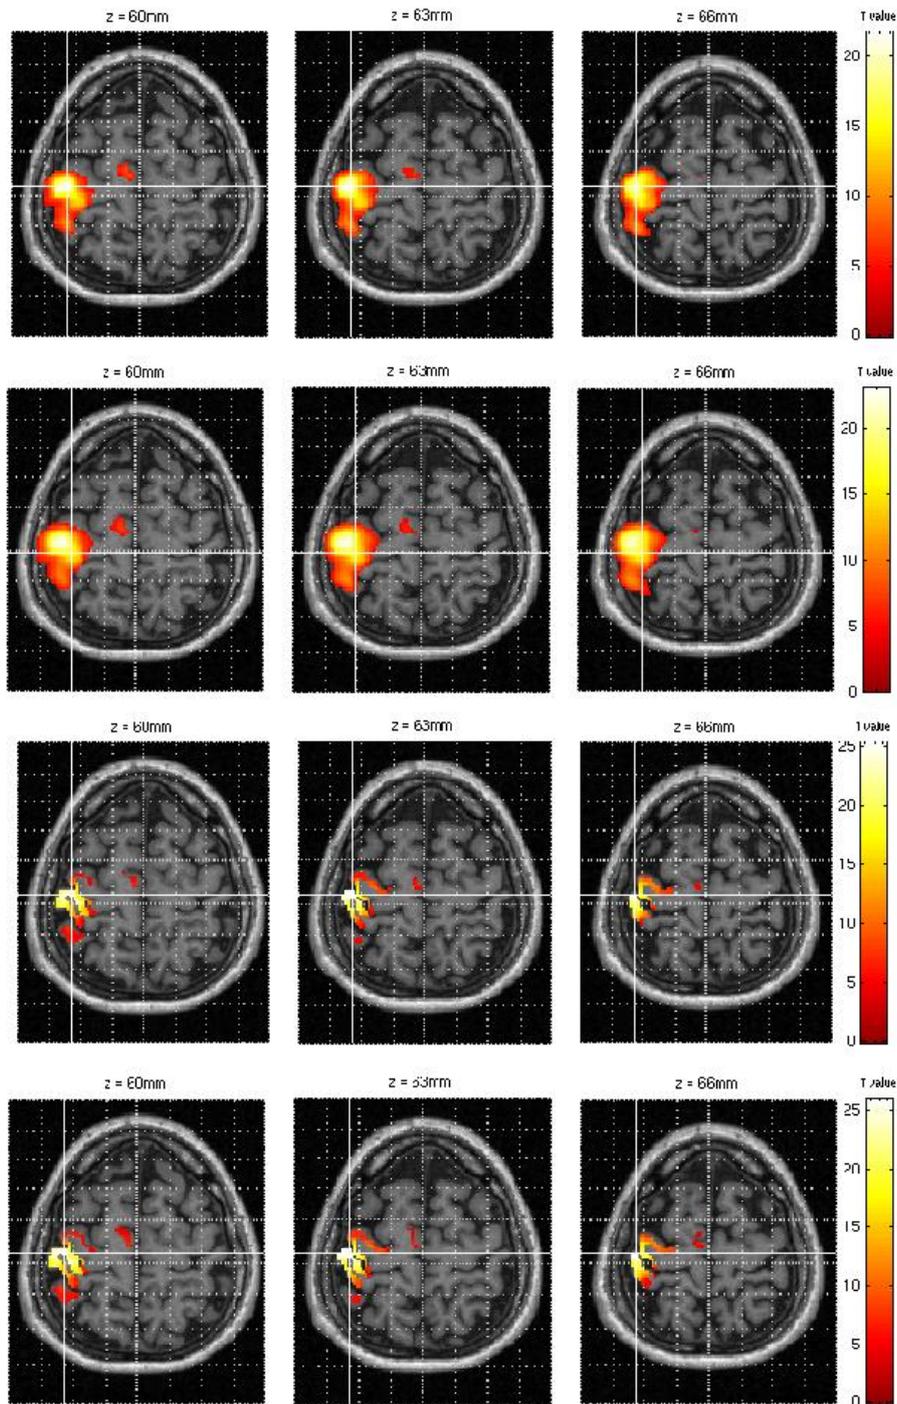


Figure 6.4: t-maps at $z=63$ mm for 4 different analyses of subject 4, thresholded at $p(t) < 0.01$ (corrected). (first row) SPM [6 6 9] mm FWHM, (second row) SPM [8 8 12] mm FWHM, (third row) AIBF [2 2 3] mm FWHM and 3 mm/3mm separation/width, (bottom row) AIBF [4 4 6] mm FWHM and 3 mm/3mm separation/width

6.2 Positron emission tomography

This section deals with the application of AIBF to PET data. The method can be applied to single subject PET data as shown in (Kiebel and Friston, 2000). In this section, two group data sets consisting of stereotactically normalized images from ten and five subjects are analyzed.

6.2.1 Multiple subject analysis

In the first analysis, it is shown that one can increase the size of local basis functions without compromising sensitivity. In contrast, a similar increase in the width of an isotropic Gaussian kernel in a conventional analysis is precluded by the matched filter theorem which requires the convolution kernel to match the spatial distribution of the signal. Furthermore, a three-dimensional smoothing filter degrades spatial resolution and confounds areas which are close in image space but quite remote on the cortical surface.

Acquisition and design

The data were obtained from five subjects scanned 12 times (every eight minutes) while performing one of two verbal tasks. Scans were obtained with a CTI PET camera (model 953B CTI Knoxville, TN USA). ^{15}O was administered intravenously as radiolabelled water, infused over two minutes. Total counts per voxel during the buildup phase of radioactivity served as an estimate of regional blood flow. Subjects performed two tasks in alternation. One task involved repeating a letter presented aurally at one per two seconds (word shadowing). The other was a paced verbal fluency task, where the subjects responded with a word that began with the letter presented (intrinsic word generation).

Analysis

For the first set of data, the VOI encompassed roughly 40,000 mm^2 of the flattened canonical surface of the left hemisphere, in particular most of the grey matter of the lateral frontal and parietal lobe. SPM99 was applied to smoothed data, with four different widths of an isotropic Gaussian kernel (16, 24, 32 and 40 mm FWHM). Four different Gaussian AIBF sets with separation/FWHM 8/10, 12/16, 16/22, 24/40 [mm] were applied with an L_I of [8 8 8] mm FWHM and an L_E of [16 16 16] mm FWHM. All AIBF results were projected into A -space, i.e. the effects due to $L = L_E L_I$ were discarded.

Results

In Fig. 6.5, one can see the maximum intensity projection (MIP) of the SPM based on smoothed data (16 mm FWHM). On the right hand side of the figure, the design matrix and the contrast are shown. This temporal model was used for all subsequent AIBF and SPM analyses.

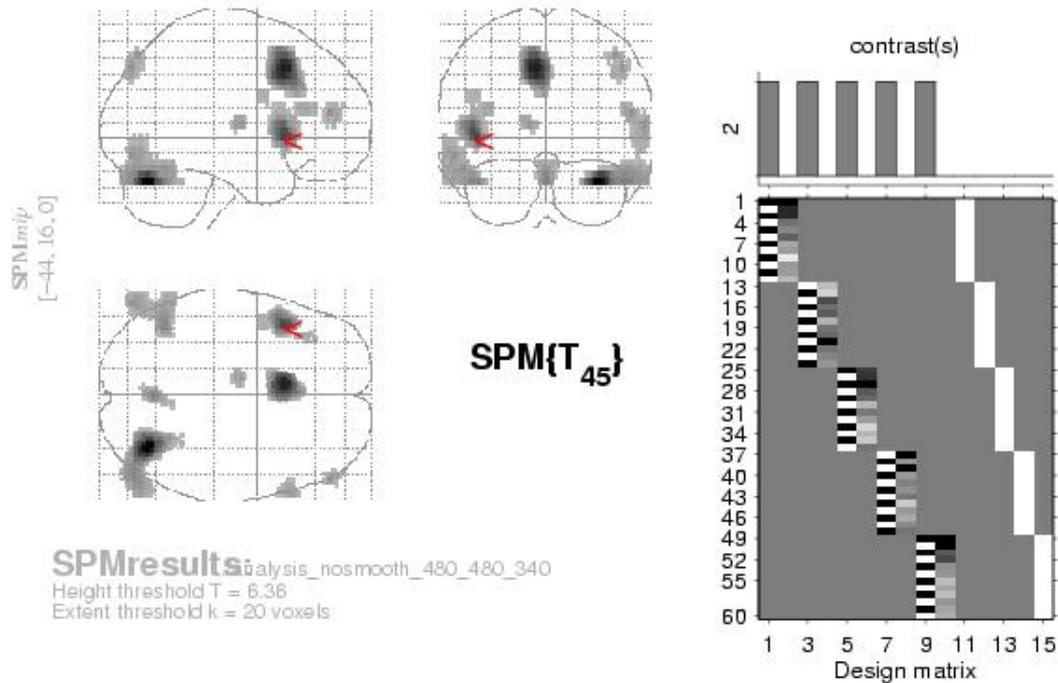


Figure 6.5: SPM analysis. Left: Maximum intensity projection (MIP) of the t-map. The red arrow marks the activation detected by all analyses. Right: design matrix and contrast used in the statistical analysis

Tables 6.3 and 6.4 list the minimum p-value for each analysis and its location in normalized anatomical space.

Discussion

In this section, AIBF and SPM were applied with large smoothing filters. Filter widths of 32 and 40 mm FWHM are usually not used for functional brain imaging analyses. However, an important point can be made by these analyses that can be generalized to other more conventional situations.

By comparing tables 6.3 and 6.4, one observes that the minimum p-value of the SPM analysis increases with increasing filter width from $3.4e^{-8}$ to $3.6e^{-5}$. An increase can

Filter FWHM [mm]	p-value	location
16	$3.4e^{-8}$	-50/20/-6
24	$3.6e^{-7}$	-48/26/-6
32	$5.0e^{-7}$	-50/26/-6
40	$3.6e^{-5}$	-52/24/-10

Table 6.3: Verbal fluency study: Minimum corrected p-values computed by SPM and their locations in voxel-space for different filter widths

Parameters [mm]	p-value	location
distance/width		
8/10	$1.2e^{-10}$	-52/20/-2
12/16	$9.7e^{-10}$	-50/20/-2
16/22	$3.0e^{-9}$	-52/20/-2
24/40	$2.5e^{-9}$	-46/32/-2

Table 6.4: Verbal fluency study: Minimum p-values computed by AIBF and their locations in voxel-space for different basis function widths

also be seen for the AIBF analyses, but only by one order of magnitude ($3.4e^{-10} \leq p \leq 2.5e^{-9}$). The estimated smoothness of the AIBF and SPM analyses, where the largest filters were used, are very similar at around 40 mm FWHM. The increase of the minimum p-value returned by SPM can be attributed to the wide smoothing kernel employed that induces partial volume effects in relation to the activation.

One feature of AIBF is that the user is given the choice between increasing smoothness in image space (via a three-dimensional stationary isotropic smoothing kernel) and increasing the smoothness on the cortical surface (by increasing the width of local basis functions). As demonstrated, the smoothing implemented by a combination of a moderate isotropic smoothing (16 mm FWHM) and a high degree of smoothness on the surface, provides a more sensitive analysis than using an isotropic smoothing of 40 mm FWHM. One could imagine the spatial model implemented by AIBF as a broad three-dimensional ribbon (the convolved basis functions) following the cortical surface, whereas the isotropic filter is an ellipsoid aligned with the image axes. The aim when specifying a smoothing width is to choose a filter size that allows a linear combination of basis functions to pick up the activation. Ideally, activations are close to the cortical sheet, but in cases when this is not true or when the spatial normalization does not find a transformation that registers the individuals' cortical surfaces and the canonical surface, one should apply an initial isotropic smoothing. This decreases resolution, but increases the robustness of the method.

Convolving the functional images with a very wide isotropic kernel decreases image resolution and can lead to the situation where two or more neighbouring sources in voxel-space cannot be resolved in the t-map. The same thing happens along the grey matter surface when one uses extremely wide basis functions, but the fundamental difference is that an isotropic kernel decreases resolution in Euclidean space, whereas

AIBF smoothing decreases resolution only along the cortical surface.

Chapter 7

Discussion

In this thesis, it has been shown that anatomically informed basis functions offer a new kind of functional data analysis. To the author's knowledge, AIBF is the first attempt to explicitly include neuroanatomical knowledge into the analysis of functional imaging data like fMRI and PET. This is in the sense that AIBF treats the analysis as a source reconstruction, as opposed to an image reconstruction problem. While conventional analyses pre-smooth the data with an isotropic spatial filter, which is not informed about the underlying neuroanatomy, AIBF effectively implement a spatially variable filter. This spatial modelling of experimentally induced effects is more efficient for brain responses distributed along the cortical surface than conventional voxel-based methods. This efficiency depends on the user-defined AIBF parameters that can be used to enhance sensitivity to the underlying activation, or increase the anatomical resolution.

7.1 Applications

There are two important applications of AIBF. The first is a high-resolution characterization of functional data. Within the AIBF framework one attempts to model the neuroanatomical structures, the neurophysiological processes and the physical measurement process that provide useful constraints on the data. This approach affords high spatial resolution in the statistical results and improves localization of the activation sources with respect to the underlying neuroanatomy. The second application is the analysis of multi-subject data. Here one fully exploits the AIBF property that smooth activations along the cortical surface are detected more efficiently than when using a spatially invariant filter kernel. Although excellent spatial normalization algorithms are available, the anatomical variance, i.e. the spatial variance in the location of a functionally specialized area with respect to macroscopic structural landmarks, makes it difficult to map individual activations to one location. Without further information, this mapping problem can only be solved by applying a spatial lowpass filter to the functional data, such that evidence for a common activation is found at the cost of resolution. AIBF can be used in combination with an isotropic spatial filter, to smooth the

data along the cortical surface. In this way, the smoothing only smears activations close to each other on the cortical surface, but leaves activations that may be neighbours in image space well resolved. Both applications are now discussed more thoroughly.

7.1.1 High-resolution modelling

An important aspect of fMRI is that the signal changes are not confined to the grey matter, but at 1.5 Tesla also originate from locations within the CSF. The exact localization of activation is not important if one is interested in a rather macroscopic characterization of activation with respect to underlying anatomical structures. However, in the context of applications like establishing detection of somatotopical or tonotopical maps, exact spatial localization can be important. Activation can be expressed, at low field strength, millimeters downstream in draining veins. This signal *dislocation* depends on the vasculature of the brain area and the scanning parameters. In this situation, one way of tackling this localization problem is to combine measured BOLD effects with prior knowledge about the individual's neuroanatomy and the physical measurement process.

In the five fMRI data sets analyzed (sec. 6.1) such *downwashing* of BOLD effect are likely be present to a greater or lesser degree. In some subjects, the most significant activation could be identified as arising from CSF. In all subjects, the resolution of $1.8 \times 1.8 \times 3 \text{ mm}^3$ was not sufficient to prevent partial volume effects in voxels at the grey matter/CSF interface. It was therefore not possible to make a statement about the exact origin of the inferred activations. However, analyzing data with AIBF models provides a more anatomically informed picture by enhancing coherent responses within the cortical surface while attenuating other components. One might anticipate that a combination of high-resolution sequences and improved spatial or spatiotemporal models of individual neuroanatomy will improve the effective anatomical resolution of fMRI measurements, given an AIBF framework.

In practice, modelling at this high level of resolution is a rather demanding task, because a precise model encompassing the distribution of a neuronal response must be known or be estimated. This model is then augmented with a model that predicts how the neuronal response is measured by MRI.

As an example, a first step towards high-resolution models is to explicitly model sources in CSF, i.e. the draining veins. Given that one can obtain venograms with high-resolution, the location of veins and their connectivity in voxel-space can be established. This would allow specification of smooth one-dimensional basis functions embedded in three-dimensional space that was confined to the veins. One assumption e.g. might be that activation at a neuronal site can cause extended activation in a vein draining this site. Knowing the orientation of a vein at any location, it should be possible to construct a forward model of how activation along this vein would be measured by any fMRI sequence. After fitting, responses observed in a given image could be attributed to their neuronal sources. In this way the signal draining veins can be used to increase sensitivity without mis-attribution of the source. Once the sources are estimated the veins can be removed from the forward model so that they are discounted

in the projection back to voxel-space. This is an extension of the approach discussed in chapter 5 where the PSF was removed from the forward model to effect a *deconvolution* (i.e. projection into A -space).

7.1.2 Group studies and sensitivity

As demonstrated in Ch. 6, group studies can be performed with AIBF to increase sensitivity. The spatial model is derived from a single surface, which is called a *canonical* surface to stress the assumption that this stereotactically normalized surface is a valid representation of all the individuals' surfaces.

In the first fMRI multisubject study, all assumptions for application of AIBF were met. The central sulci were properly aligned with a maximal translational error of roughly 1 - 2 mm. The individual activations consisted of an elongated activation cluster along the surface and anatomical variability was presumably small and confined to the surface. The result, an increase in sensitivity and better effective resolution compared to a conventional analysis was compelling. Only a small Gaussian isotropic filter was required to allow for misregistration of individual activations. The second example, a PET study, was used to show the difference between conventional smoothing and AIBF. The point was made that the effective smoothing with AIBF follows the surface such that less non-signal related intensity confounds the parameter estimation as compared to a conventional analysis based on an isotropic invariant filter.

The key point about AIBF is that, for each image location, it effectively implements a filter kernel that resembles the shape of the signal given the underlying neuroanatomy. As a result, the difference in sensitivity between AIBF and SPM depends on the closer approximation to the activation's shape. If the signal width is large, AIBF will find a better filter than a voxel-based method. However, if the signal width is at the sub-voxel level, no spatial smoothing provides the best model. Generally, the wider the signal on the cortical sheet the more sensitive AIBF is in comparison to a conventional method. This is because, for wide signals on the cortical sheet, the most efficient AIBF smoothing ensues, if the resulting smoothing kernel at each voxel has a ribbon-like elongated shape following the sheet. These smoothing kernels can only be generated by AIBF parameters, where the width of the basis functions tangential to the cortical sheet is large compared to the width in the normal direction. If this is not the case (i.e. the width of the kernel in normal direction is larger than in tangential direction of the surface), the AIBF smoothing kernel at each voxel would resemble the shape of an isotropic invariant smoothing kernel and the advantage of an anatomically informed smoothing kernel would be lost.

With AIBF, the user can choose which space the fitted parameters are projected onto. The space in which the parameters are estimated is called A_L -space. In this space the basis functions have maximum smoothness because of the convolution with L_I and L_E and the corresponding statistical maps will be very smooth. This space can be used to assess group results because the convolution with L_E smears individual activations (c.f. the smoothing filter applied in conventional voxel-based analysis). The A -space, i.e. the space without the effects of convolution with L_E or L_I can be used to find the

locations of underlying activations on the cortical surface. Together with the choice of basis functions, AIBF provides a large range of possible spatial models.

The approach described in this thesis uses transformation parameters derived from the structural MRI where the normalization is based on the constrained minimization of the difference between the image and the template image, but not of a function of a surface and its template. For higher-order brain areas it does not seem to be possible to infer the localization of specific functional areas from structural macroscopical landmarks (Amunts et al., 1999). For primary areas there seems to be some relationship between structural and functional features as indicated by Amunts et al. (2000). In these cytoarchitectonic studies, the variability of macroscopic landmarks in relation to cytoarchitectonically defined functional areas were estimated. This generates *probability maps*, which encode the estimated probability of a given cytoarchitectonic area in a population at a given coordinate in Talairach-space (after affine transformations of MRI-registered post-mortem brains). These studies (Amunts et al., 2000) indicate that there is some anatomical variability of functional borders even in primary areas like Brodmann area (BA) 17, i.e. V1 (visual area 1). However, given the typical resolution of functional data (e.g. 4 - 8 mm FWHM), it appears that anatomical variability in BA 17 is low compared to the resolution of statistical results. It was also shown in (Amunts et al., 2000) that the anatomical variability of BA 18, i.e. V2 (visual area 2) is higher than in BA 17. The anatomical variability is even higher in areas like BA 44/45. These results speak to a non-stationary distribution of anatomical variability in voxel-based Talairach-space, where it is least in primary areas and increases in higher-order areas. Clearly, probability maps as assessed in (Amunts et al., 1999) would be most naturally represented on the cortical surface space, because it is in this space where probability maps are assessed. An interesting question is whether it is worthwhile applying surface-based normalization for studies in which one is interested in areas with low anatomical variability. A surface-based normalization would align individuals' surfaces such that residual anatomical variability is tangential to the cortical surface. This would make sense in primary areas, where a non-linear image-based normalization does not align surfaces properly, leaving some residual error in the normal direction. As found in the first experiment (fMRI), this error is small in the central sulcus. Presumably this is because the central sulcus is a sharp-edged feature embedded within homogeneous white matter such that an intensity-based normalization can easily align this feature across subjects. One has to keep in mind that the resolution of the analysis is related to the voxel-size, i.e. a residual surface alignment error of 1 mm would, in most studies, translate to an error of only a third or half a voxel. Given that anatomical variability within a functional area has yet to be quantified, it is questionable that a surface-based normalization would improve the sensitivity of an analysis in the primary sensorimotor areas. Other candidates for surface-based normalization would be e.g. primary auditory and visual areas. However, as long as anatomical variability in these areas remains unquantified and in the absence of comparing a non-linear intensity-based normalization to a surface-based one, no conclusion can be drawn about the relative benefits of intensity- and surface-based normalization.

7.2 Error sources and Limitations

There are some potential error sources and limitations to AIBF as it has been implemented in this thesis. None of these compromise the concept of AIBF, but the tolerance of AIBF can be rather narrow, in particular when the aim is to maximize resolution. If one increases the robustness of the model by smoothing both the data and the model with an isotropic filter, most of the error sources are less of an issue.

The first obvious potential error is the distortion in the most often used fMRI sequence, EPI. The cortical surface is reconstructed based on a structural sequence, which is undistorted. This mismatch can preclude an exact registration of surface and functional data. The principled way of solving this problem is to apply a distortion correction to the functional data.

Another error source is the coregistration between structural and functional images.

The surface reconstruction is a potentially large source of error. Although the quality of the surface reconstructions used within this thesis was sufficient to define a spatial model, there are some potential pitfalls that can compromise the anatomical model. For completeness, these are as follows: The structural sequences have to be of high-resolution (i.e. at least $1 \times 1 \times 1 \text{ mm}^3$ voxel-size). Most surface reconstructions computed by the author produced a small number of *defects*, i.e. some topology of the surface was incorrectly assessed. Most of these effects are negligible and do not have a substantial effect on the model specification or the model fit. However, care must be taken to remove large defects from the surface reconstruction before specifying an AIBF model. Another issue is the white-grey contrast in the structural T1-weighted image. Most structural sequences have only optimal contrast near to the centre of the image and less contrast away from the centre. This can effect the proper localization of the grey-white matter interface and consequently inhomogeneity correction should be applied prior to surface reconstruction. Even at a high-resolution, partial volume effects at the grey-white interface occur. In brain regions where the white matter volume enclosed between the banks of two adjacent sulci is close to the voxel-size, the white matter volume is underestimated. This can happen in the lower parts of the temporal lobes, even when using a voxel-size of $[1 \ 1 \ 1] \text{ mm}^3$.

The AIBF model is based on a flattened surface. There are two problems to deal with. The first is that the transformation from folded to flattened surface introduces geometrical distortions in the surface, i.e. distances between two points on the surface are changed. While the distribution of local basis functions defined on the flattened surface can be described by the two parameters (separation and width), this is no longer true on the folded surface, because all basis functions experience local distortions. Although it might be perceived that this is a weakness in the model specification, it should be noted that this only alters the local smoothness of the basis functions. Consequently, the fit can reflect this variation in smoothness. In other words, the local smoothness of the fit depends on the distribution of surface distortions. In this work, all brain regions were distorted at most by roughly 10%, which is negligible given that one is only interested in a smooth distribution on the cortical surface. Another potential problem is the introduction of cuts to flatten the surface. Clearly, if a cut is made

through an activation location, the ensuing spatial model is no longer valid.

In practice, one limitation of the current implementation is that subcortical structures are not modelled, e.g. thalamic or cerebellar activation is not detectable. This can be overcome by specifying basis functions for these structures.

Storing the basis functions and solving the equations needs a large amount of memory. In the current implementation, this problem is overcome by only analyzing a volume of interest, where the maximum possible size of the volume depends on the size of the examined voxels, the number of basis functions and their width on the cortical surface. Another approach is to apply matrix inversion algorithms that exploit the sparse nature of most AIBF models.

Chapter 8

Zusammenfassung

Ein Teilgebiet der Neurowissenschaften ist die Erforschung des menschlichen Gehirns mittels makroskopisch bildgebender Verfahren wie funktionelle Magnetresonanztomographie (fMRT) oder Positronenemissionstomographie (PET). Beide Verfahren messen u.a. durch neuronale Aktivität induzierte metabolische Prozesse und bieten eine räumliche Auflösung im Millimeter-Bereich an. Die Kenntnis der zeitlichen Abfolge experimenteller Bedingungen erlaubt es, durch das Experiment verursachte Effekte zu detektieren und zu interpretieren. Fast alle für die Analyse von Gehirnbilddaten entwickelten Methoden haben ihren Schwerpunkt auf die Modellierung zeitlicher Effekte gelegt, während vergleichsweise wenig Entwicklungsarbeit in die angemessene Charakterisierung der räumlichen Information geflossen ist. Das üblicherweise verwendete räumliche Modell der Aktivierungen benutzt nicht die (verfügbaren) neuroanatomischen Informationen über den Probanden oder eine Gruppe von Probanden, sondern die implizite Annahmen, daß (i) an jeder Stelle innerhalb des Gehirns Aktivierung gleichwahrscheinlich ist und (ii) die räumliche Verteilung der Aktivierung am besten durch einen dreidimensionalen, invarianten und mit den Bildachsen ausgerichteten Tiefpaßfilter beschrieben wird. Beide Annahmen sind umso *inkorrekt*, desto höher die räumliche Auflösung des Meßapparats ist.

In dieser Arbeit wird die Methode der anatomisch informierten Basisfunktionen (AIBF) beschrieben, mit der existierende neuroanatomische Informationen in die Analyse funktioneller Gehirnbilddaten eingebracht werden können. Eine offensichtliche Wahl anatomischer Information ist die Oberfläche, die in einem individuellen Gehirn die graue Substanz approximiert. Für die Oberflächenextraktion wurde das Software-Paket Brainvoyager 3.8 (Rainer Göbel, Brain Innovations) benutzt, das auf einem hochauflösenden ($1 \times 1 \times 1 \text{ mm}^3$) T_1 -gewichteten MRT-Bild arbeitet. Gegeben die zwei Annahmen, daß (i) Aktivierungen in oder nahe der grauen Substanz lokalisiert sind und (ii) Aktivierungen räumlich glatt in Richtung der Oberfläche der grauen Substanz verlaufen, kann ein räumliches lineares Modell formuliert werden. Dieses Modell erklärt für jedes Bild einer funktionellen Bildserie die Verteilung der Bildintensitäten durch eine Linearkombination von räumlich niederfrequenten Basisfunktionen, die entlang der Oberfläche der grauen Substanz definiert sind. Die Schätzung der zugehörigen Modellparameter implementiert eine Reduktion von zehntausenden von Voxelintensitäten auf wenige

hundert Parameter. Die Rückprojektion dieser Parameterschätzungen in den Voxelraum ist äquivalent zu der Schätzung eines variablen, räumlichen Filters, wobei die Form des lokalen Filters durch den Verlauf der gyralen Struktur der grauen Substanz bestimmt wird. Die zeitliche Analyse der geschätzten oder rückprojizierten Parameter kann dann mit klassischen Methoden, z.B. statistical parametric mapping (SPM) durchgeführt werden. Der wesentliche Verarbeitungsschritt liegt in der Reduktion der funktionellen Bilder auf die geschätzte räumlich-zeitliche Parametermatrix. Hierbei werden nur jene Komponenten extrahiert, die räumlich niederfrequent entlang der kortikalen Oberfläche verlaufen. Falls die obigen beiden Annahmen korrekt sind, dann wird diese anatomisch informierte Filtermethode einen experimentell induzierten Effekt mit hoher Sensitivität detektieren. Es wird ebenfalls erwartet, daß die effektive Auflösung zwischen Aktivierungsquellen im Vergleich zu konventionellen Filtermethoden steigt, weil der anatomisch informierte Filter nur entlang der Oberfläche glättet und im Bildraum benachbarte Quellen (die aber auf der kortikalen Oberfläche weit entfernt liegen) unterscheidbar läßt.

Mittels simulierter Daten wurde gezeigt, daß AIBF bei Erfüllung der Annahmen tatsächlich sensitiver ist als konventionelle auf räumlich invarianten Filtern basierende Methoden.

Bei Anwendung von AIBF auf reale fMRT- und PET-Daten werden drei verschiedene Anwendungen demonstriert. Die erste ist eine exakte räumlich hochauflösende Modellierung der funktionellen Daten eines einzelnen Probanden in Bezug auf bekannte physiologische, physikalische und sonstige Effekte, so daß es möglich ist, ein räumliches Modell zu generieren, mit dem der experimentell induzierte BOLD-Effekt sehr genau in einem Probanden lokalisiert werden kann. Es wurde gezeigt, daß die Benutzung der kortikalen Oberfläche als anatomisches Wissen die Generierung eines Modells erlaubt, mit dem prinzipiell anatomisch spezifischere Aussagen getroffen werden können. In den vorliegenden realen fMRT-Daten war es z.B. möglich, tendentiell zwischen Aktivierung innerhalb von zerebraler Flüssigkeit und grauer Substanz zu unterscheiden, so daß der Ursprung von Aktivierungen besser lokalisiert werden kann.

Die zweite Anwendung hat nicht die Verbesserung der anatomischen Präzision zum Ziel, sondern die Sensitivität und Robustheit der AIBF-Methode. Hierzu wurde anhand von fMRT-Daten gezeigt, daß ein zusätzlicher auf das anatomisch informierte Modell und die Daten angewandter räumlich invarianter Tiefpaß-Filter anatomisch nicht korrekt modellierte Aktivierungen auf Kosten der Auflösung zu detektieren vermag.

Die dritte Anwendung ist vielleicht die interessanteste: Die Methode wird zur Analyse von Gruppendaten eingesetzt. Auf den ersten Blick macht das nicht viel Sinn, weil die sogenannte *anatomische Variabilität* (die unterschiedliche Lokalisierung des Ortes einer Funktion in räumlicher Relation zu makroskopisch markanten strukturellen Punkten gemessen über Probanden) die anatomisch exakte Modellierung der Funktion eines stereotaktisch normalisierten Areals verhindert. In dieser Situation wird in Gruppenstudien bei konventionellen Methoden ein räumlicher, relativ weiter, dreidimensionaler, invarianter Tiefpaßfilter benutzt um die durch den Einfluß der anatomischen Variabilität verteilten individuellen Aktivierungen zu verschmieren und die Gruppenaktivierung dann mit hoher Sensitivität im Überlappungsgebiet der individuellen

Aktivierungen zu detektieren. Die hier vorgeschlagene und implementierte Methode ist ähnlich, nur daß der räumliche Tiefpaßfilter aus zwei Komponenten besteht. Die erste modelliert mittels Basisfunktionen die Verteilung der individuellen Aktivierungen entlang der kortikalen Oberfläche und die zweite erklärt durch einen dreidimensionalen, invarianten Filter u.a. alle Effekte in die Normalenrichtung der Oberfläche. Nach der Parameterschätzung wird bei der Rückprojektion allerdings in den Raum projiziert, der nur die erste Komponente entlang der Oberfläche modelliert. Hierdurch wird eine anatomisch informierte Entfaltung implementiert, so daß die Effekte des invarianten Filters aus den statistischen Ergebnissen entfernt werden. Diese Anwendung erlaubt Gruppenstudien mit erhöhter Sensitivität und einer höheren anatomischen Lokalisierung.

Acknowledgments

Ein großer Teil meiner Dankbarkeit gebührt zwei Männern: Cornelius Weiller und Karl Friston. Ohne diese beiden wäre meine Dissertation vermutlich nicht zustande gekommen. Cornelius hat mir in Essen und Jena in einer neurologischen Universitätsklinik die vergleichsweise ruhige Arbeitsumgebung und immer wohlwollende Unterstützung angedeihen lassen, die ich für diese Arbeit brauchte. Auch hat Cornelius in dieser Zeit die effektive Zusammenarbeit mit Karl ermöglicht. Karl war/ist der Betreuer meiner Arbeit. Durch seinen initialen motivierenden Anschlag und konsistente, konstruktive und geduldige Unterstützung konnte ich die Arbeit vollenden.

Mein weiterer Dank für sehr hilfreiche Diskussionen und viel Unterstützung gehört Lutz Jäncke, Rainer Göbel, Volkmar Glauche, Christian Gaser, Bruce Fischl, Andrew Holmes, Jean-Baptiste Poline, John Ashburner, Oliver Josephs, Dave McGonigle, Ralf Deichmann, Christophe Phillips, Chloe Hutton, Will Penny und Christian Büchel.

Curriculum vitae

Persönliche Daten

Name	Stefan Johannes Kiebel
Anschrift	57 Fairfax Road London N8 0NH United Kingdom
Geb.	10. Juni 1968 in Wadern/Saar
Nationalität	Deutsch

Schule

1974 - 1978	Grundschule an der Maarbrücke, Bochum
1978 - 1987	Gymnasium am Ostring, Bochum

Wehrdienst

1987 - 1988	Wehrdienst in Achim bei Bremen
-------------	--------------------------------

Studium

1988-1994	Studium der Informatik, Universität Dortmund
1993/1994	Diplomarbeit im Institut für Neuroinformatik, Ruhr-Universität Bochum

Berufstätigkeit

Oktober 1994 - Mai 1996

Mitglied der funktionellen Bildverarbeitungsgruppe
(Leiter: Cornelius Weiller) und Systemadministrator
in der Klinik für Neurologie
in der Universitätsklinik Essen

Mai 1996 - Oktober 1999

Mitglied der funktionellen Bildverarbeitungsgruppe
(Leiter: Cornelius Weiller) und Systemadministrator
in der Klinik für Neurologie
in der Universitätsklinik Jena

Seit Oktober 1999

Mitglied der Methodengruppe
(Leiter: Karl Friston) des
Functional Imaging Laboratory,
Wellcome Dept. of Cognitive Neurology,
Queen Square, London, UK

Bibliography

- Aguirre, G., Zarahn, E., and D’Esposito, M. 1997. Empirical analyses of BOLD fMRI statistics. *NeuroImage*, 5:199–212.
- Aguirre, G., Zarahn, E., and d’Esposito, M. 1998. The variability of human, BOLD hemodynamic responses. *Neuroimage*, 8:360–369.
- Amunts, K., Malikovic, A., Mohlberg, H., Schormann, T., and Zilles, K. 2000. Brodmann’s area 17 and 18 brought into stereotaxic space — where and how variable. *Neuroimage*, 11:66–84.
- Amunts, K., Schleicher, A., Burgel, U., Mohlberg, H., Uylings, H., and Zilles, K. 1999. Broca’s region revisited: cytoarchitecture and intersubject variability. *J of Comp Neurol*, 20:319–341.
- Andersson, J., Vagnhammar, B., and Schneider, H. 1995. Accurate attenuation correction despite movement during PET imaging. *J. of Nuclear Medicine*, 36:670–678.
- Ashburner, J. and Friston, K. 1997. Multimodal image coregistration and partitioning — a unified framework. *Neuroimage*, 6:209–217.
- Ashburner, J. and Friston, K. 1999. Nonlinear spatial normalization using basis functions. *Human Brain Mapping*, 7:254–266.
- Ashburner, J. and Friston, K. 2000. Voxel-based morphometry — the methods. *Neuroimage*, 11:805–821.
- Ashburner, J., Hutton, C., Frackowiak, R., Johnsrude, I., Price, C., and Friston, K. 1998. Identifying global anatomical differences: deformation-based morphometry. *Human Brain Mapping*, 6:348–357.
- Carman, G., Drury, H., and van Essen, D. 1995. Computational methods for reconstructing and unfolding the cerebral cortex. *Cerebral Cortex*, 5:506–517.
- Chatfield, C. and Collins, A. 1980. *Introduction to Multivariate Analysis*. Chapman and Hall, London.
- Cohen, M. S. 1997. Parametric analysis of fMRI data using linear systems methods. *NeuroImage*, 6:93–103.
- Collins, D., Neelin, P., Peters, T., and Evans, A. 1994. Automatic 3D intersubject registration of MR volumetric data in standardized talairach space. *J Comput Assist Tomogr*, 18:192–205.
- Dale, A. M., Fischl, B., and Sereno, M. I. 1999. Cortical surface-based analysis, I. segmentation and surface reconstruction. *Neuroimage*, 9:179–194.
- Dale, A. M. and Sereno, M. I. 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach. *Journal of Cognitive Neuroscience*, 5:162–176.
- Deichmann, R. and Turner, R. 2000. Improved t1-weighted imaging. *Neuroimage*.

- Engel, S., Glover, G., and Wandell, B. 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, 7:181–192.
- Engl, H. W. and Grever, W. 1994. Using the L-curve for determining optimal regularization parameters. *Numerische Mathematik*, 69:25–31.
- Fischl, B. and Dale, A. M. 2000. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proc Natl Acad Sci USA*, 97:11050–5.
- Fischl, B., Sereno, M., Tootell, R., and Dale, A. 1999a. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping*, 8:272–284.
- Fischl, B., Sereno, M. I., and Dale, A. M. 1999b. Cortical surface-based analysis, II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9:195–207.
- Frackowiak, R., Friston, K., D. Frith, C., Dolan, R., and Mazziotta, J. C. 1997. *Human Brain Function*. Academic Press.
- Friston, K., Ashburner, J., Frith, C., Poline, J.-B., Heather, J., and Frackowiak, R. 1995a. Spatial registration and normalization of images. *Human Brain Mapping*, 2:165–189.
- Friston, K., Frackowiak, K. W. R., Mazziotta, J., and Evans, A. 1994a. Assessing the significance of focal activations using their spatial extent. *Human Brain Mapping*, 1:214–220.
- Friston, K., Holmes, A., Poline, J.-B., Price, C., and Frith, C. 1996. Detecting activations in PET and fMRI: Levels of inference and power. *Neuroimage*, 4:223–235.
- Friston, K., Jezzard, P., and Turner, R. 1994b. Analysis of functional MRI time-series. *Human Brain Mapping*, 1:153–171.
- Friston, K., Josephs, O., Zarah, E., Holmes, A., Rouquette, S., and Poline, J.-B. 2000. To smooth or not to smooth? Bias and efficiency in fMRI time-series analysis. *Neuroimage*, 12:196–208.
- Friston, K. J., Frith, C. D., Frackowiak, R. S. J., and Turner, R. 1995b. Characterizing dynamic brain responses with fMRI. *NeuroImage*, 2:166–172.
- Gaser, C., Volz, H., Kiebel, S., Riehemann, S., and Sauer, H. 1999. Detecting structural changes in whole brain based on nonlinear deformations-application to schizophrenia research. *Neuroimage*, 10:107–113.
- Gonzalez, R. C. and Wintz, P. 1987. *Digital Image Processing*, chapter 3, page 121. Addison-Wesley Publishing Company, Inc.
- Haacke, E. M., Brown, R. W., Thompson, M. R., and Venkatesan, R. 1999. *Magnetic Resonance Imaging*. John Wiley and Sons, Inc.
- Hansen, P. C. 1992. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34:561–580.
- Harville, D. A. 1974. Bayesian inference for variance components using only error contrasts. *Biometrika*, 61:383–385.
- Harville, D. A. 1977. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association*, 72:320–338.
- Holmes, A. 1994. *Statistical Issues in Functional Brain Mapping*. PhD thesis, University of Glasgow.
- Jezzard, P. and Balaban, R. S. 1995. Correction for geometric distortion in echo planar images from B_0 field variations. *Magnetic Resonance in Medicine*, 34:65–73.
- Jezzard, P. and Clare, S. 1999. Sources of distortion in functional MRI data. *Human*

- Brain Mapping*, 8:80–85.
- Jones, S. E., Buchbinder, B. R., and Aharon, I. 2000. Three-dimensional mapping of cortical thickness using Laplace’s equation. *Human Brain Mapping*, 11:12–32.
- Kastrup, A., Kruger, G., Glover, G., Neumann-Haefelin, T., and Moseley, M. 1999. Regional variability of cerebral blood oxygenation response to hypercapnia. *Neuroimage*, 10:675–681.
- Kiebel, S. J., Ashburner, J., Poline, J.-B., and Friston, K. J. 1997. MRI and PET coregistration — a cross validation of statistical parametric mapping and automated image registration. *Neuroimage*, 5:271–279.
- Kiebel, S. J. and Friston, K. J. 2000. Characterization of PET data using anatomically informed basis functions. *Neuroimage*, 11(5):463.
- Kimmel, R. and Sethian, J. 1998. Computing geodesic paths on manifolds. *Proc. Natl. Acad. Sci. USA*, 95:8431–8435.
- Kruggel, F., von Cramon, D. Y., and Descombes, X. 1999. Comparison of filtering methods for fMRI datasets. *Neuroimage*, 10:530–543.
- Lange, N. and Zeger, S. L. 1997. Non-linear fourier time series analysis for human brain mapping by functional magnetic resonance imaging. *Applied Statistics, Journal of the Royal Statistical Society, Series C*, 46(1):1–29.
- MacDonald, D., Kabani, N., Avis, D., and Evans, A. 2000. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. *NeuroImage*, 12:340–356.
- Press, W. H., Teukolsky, S. A., and Vetterling, W. T. 1992. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press.
- Rao, C. R. and Toutenburg, H. 1995. *Linear Models, Least Squares and Alternatives*. Springer Series in Statistics.
- Schwartz, E. and Merker, B. 1986. Computer-aided neuroanatomy: Differential geometry of cortical surfaces and an optimal flattening algorithm. *IEEE Comp. Graph. Appl.*, 6:36–44.
- Schwartz, E., Shaw, A., and Wolfson, E. 1989. A numerical solution to the generalized mapmaker’s problem: Flattening nonconvex polyhedral surfaces. *IEEE Transactions on Pat Anal Mach Intel*, 11:1005–1008.
- Sereno, M., Dale, A., Reppas, J., Kwong, K., Belliveau, J., Brady, T., Rosen, B., and Tootell, R. 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, 268:889–893.
- Talairach, J. and Tournoux, P. 1988. *Coplanar stereotaxic atlas of the human brain*. New York, Thieme Medical.
- Thompson, P., MacDonald, D., Mega, M., Holmes, C., Evans, A., and Toga, A. 1997. Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces. *Journal of Computer Assisted Tomography*, 21:567–581.
- Tootell, R., Mendola, J., Hadjikhani, N., Ledden, P., Liu, A., Reppas, J., Sereno, M., and Dale, A. 1997. Functional analysis of V3a and related areas in human visual cortex. *Journal of Neuroscience*, 17:7060–7078.
- van Essen, D. and Drury, H. 1997. Structural and functional analyses of human cerebral cortex using a surface-based atlas. *Journal of Neuroscience*, 17:7079–7102.
- van Essen, D. and Zeki, S. 1978. The topographic organization of rhesus monkey prestriate cortex. *Journal of Physiology*, 277:193–226.
- Woods, R., Cherry, S., and Mazziotta, J. 1992. Rapid automated algorithm for aligning and reslicing PET images. *J. Comput. Assist. Tomography*, 16:620–633.

- Woods, R., Grafton, S., Watson, J., Sicotte, N., and Mazziotta, J. 1998. Automated image registration: II. intersubject validation of linear and nonlinear models. *J comput Assist Tomogr*, 22:153–165.
- Worsley, K., Marrett, S., Neelin, P., Vandal, A., Friston, K., and Evans, A. 1995. Searching scale space for activation in PET images. *Human Brain Mapping*, 4:74–90.
- Worsley, K. J., Andermann, M., Koulis, T., MacDonald, D., and Evans, A. 1999. Detecting changes in non-isotropic images. *Human Brain Mapping*, 8:98–101.
- Worsley, K. J. and Friston, K. J. 1995. Analysis of fMRI time-series revisited — again. *Neuroimage*, 2:173–181.
- Worsley, K. J., Marrett, S., Neelin, P., Vandal, A., Friston, K., and Evans, A. 1996. A unified statistical approach for determining significant signals in images of cerebral activation. *Human Brain Mapping*, 4:58–73.
- Yousry, T., U.D. Schmid, H. A., Schmidt, D., Peraud, A., Buettner, A., and Winkler, P. 1997. Localization of the motor hand area to a knob on the precentral gyrus. *Brain*, 120:141–157.
- Zarahn, E., Aguirre, G., and M.D’Esposito 1997. Empirical analyses of BOLD fMRI statistics. *Neuroimage*, 5:179–197.