Chapter 1

Introduction

Contents

1.1 Motivation and Aims

The initial motivation for this work was to develop improved methods of image registration for functional imaging. Much of it has now been incorporated into the SPM99 package, and is used by several hundred researchers around the world for analysing functional imaging data. The second motivation was to facilitate the development of methods for studying brain shape among different populations. Recently, the term *computational neuroanatomy* has been coined for this area of research. These areas are now reviewed in more detail.

1.1.1 Functional Imaging

The principle behind detecting activations using functional imaging methods such as *Positron Emission Tomography* (PET) or *functional Magnetic Resonance Imaging* (fMRI) is essentially a voxel by voxel t-test on a series of images acquired under different conditions (Friston *et al.*, 1995d; Worsley & Friston, 1995). This analysis results in a statistical parametric map (SPM) showing significant differences in cerebral blood flow that are explained according to the different conditions experienced by the subject (or subjects) in the scanner.

The first modalities used for this type of study were PET and SPECT (*Single Photon Emission Computed Tomography*). Recent advances in MR methods, the large number of existing scanners, coupled with the relatively high cost of producing radioactive tracers, and the invasive nature of PET and SPECT have meant that most studies currently use fMRI.

PET and SPECT methods involve generating images from the photons of radiation emitted by tracers injected into, or inhaled by, the subjects. Tracers used in PET studies emit positrons

when they decay, and include ^{11}C, ^{15}O and ^{18}F, although most studies are based on identifying regional differences in cerebral blood flow, and involve injecting the subject with ^{15}O labelled water, or the subject breathing ^{15}O labelled carbon dioxide. When a positron meets an electron, the two annihilate each other resulting in the emission of two gamma ray photons in opposite directions. By recording the paths of the gamma ray pairs, it is possible to reconstruct a three dimensional image of the tracer concentration. For PET studies, a typical protocol may involve about 12 scans, with an injection of ^{15}O labelled water prior to each of them. The subject may have one task to do for six of the scans, and a different task for the other six. More active brain regions have a higher rate of blood flow, and so receive the tracer earlier than the other areas. By imaging the brain as the radioactivity is entering, and comparing the images resulting from the different tasks, it is possible to create a picture of where the tasks have most influence over the cerebral blood flow.

The mechanisms of MRI are very different from those of PET, and rely on the nuclei of certain atoms (normally ^{1}H) absorbing and then re-emitting radio waves when in a magnetic field. The frequency of the absorbed and emitted waves depends on the strength of the magnetic field, so by varying the field over the head it is possible to record waves of different frequencies from different regions. Fourier transform methods are used to reconstruct images of where the signals emanate from. Depending on the properties of the surrounding tissue, the amplitudes of the signals will decay at different rates, so not only does MRI produce a map of the density of ^{1}H atoms, but it also says something about the environment in which the atoms are found.

Currently, the index of neuronal activity most commonly used for fMRI is the *Blood Oxygenation Level Dependent* (BOLD) contrast (Ogawa *et al.*, 1990). The assumption is that an increase in neuronal activity within a brain region results in an increase in local blood flow, leading to reduced concentrations of deoxyhæmoglobin in the blood vessels. Unlike oxyhæmoglobin, deoxyhæmoglobin has a differential magnetic susceptibility in relation to the surrounding tissue. Therefore, relative decreases in deoxyhæmoglobin concentration lead to a reduction in local field inhomogeneity and a slower decay of the MR signal, resulting in higher intensities in the images. fMRI allows whole brain images to be collected in about six seconds, giving it a much better temporal resolution than PET (as there is typically a wait of about 10 minutes between scans in order for the radiation from the previous scan to decay). This means that hundreds of fMRI volumes are often collected for each subject.

Image registration is important in many aspects of functional image analysis. In imaging neuroscience, particularly for fMRI, the signal changes due to any hæmodynamic response can be small compared to apparent signal differences that can result from subject motion, so it is important that the images are as closely aligned as possible prior to performing the statistical tests. Subject head movement in the scanner can not be completely eliminated, so motion correction needs to be performed as a preprocessing step. Most current algorithms for movement correction consider the head as a rigid object. In three dimensions, six parameters are needed to define a rigid body transformation (three translations and three rotations). The first step in the correction is image registration, which involves determining the parameter values for a rigid body transformation that optimise some criteria for matching each image with a reference image. In order to be effective, the accuracy of the registration needs to be within a fraction of a voxel. Following the registration, the images are transformed by resampling according to the determined parameters.

Motion correction is especially important for experiments where subjects may move in the scanner in a way that is correlated with the different conditions (Hajnal *et al.*, 1994). Even tiny systematic differences can result in a significant signal accumulating over numerous scans. Without suitable corrections, artifacts arising from subject movement correlated with the experimental paradigm may appear as activations. A second reason why motion correction is important is that it increases sensitivity. The t-test is based on the signal change relative to the residual variance. The residual variance is computed from the sum of squared differences between the data and the linear model to which it is fitted. Movement artifacts add to this residual variance, and so reduce the sensitivity of the test to true activations.

For studies of a single subject, sites of activation can be accurately localised by superimposing them on a high resolution structural image of the subject (typically a T1 weighted MRI). This requires registration of the functional images with the structural image. As in the case of movement correction, this is normally performed by optimising a set of parameters describing a rigid body transformation, but the matching criterion needs to be more complex because the structural and functional images normally look very different. A further use for this registration is that a more precise spatial normalisation can be achieved by computing it from a more detailed structural image. If the functional and structural images are in register, then a warp computed from the structural image can be applied to the functional images.

Sometimes it is desirable to warp images from a number of individuals into roughly the same standard space to allow signal averaging across subjects. This procedure is known as spatial normalisation. Because different people may have different strategies for performing tasks in the scanner, spatial normalisation of the images is useful for determining what happens generically over individuals. A further advantage of using spatially normalised images is that activation sites can be reported according to their Euclidian co-ordinates within a standard space (Fox, 1995). The most commonly adopted co-ordinate system within the brain imaging community is that described by Talairach & Tournoux (1988), although new standards are now emerging that are based on digital atlases (Evans *et al.*, 1993; Evans *et al.*, 1994; Mazziotta *et al.*, 1995).

1.1.2 Computational Neuroanatomy

A large number of approaches for characterising differences in the shape and neuroanatomical configuration of different brains have recently emerged due to improved resolution of anatomical human brain scans and the development of new sophisticated image processing techniques.

One of the simplest morphometric approaches involves identifying shape changes within single subjects by subtracting coregistered images acquired at different times. The changes could be because of a number of different reasons, but most are related to pathology. Because the scans are of the same subject, the first step for this kind of analysis involves registering the images together by a rigid body transformation.

Other approaches require the images of multiple subjects to be registered together by some form of spatial normalisation. The primary result of spatially normalising a series of images is that they all conform to the same stereotactic space, enabling region-by-region comparisons to be performed. A second result is a series of deformation fields that describe the spatial transformations required to match the different shaped brains to the same template. Encoded within each deformation field is information about the individual image shapes, which can be

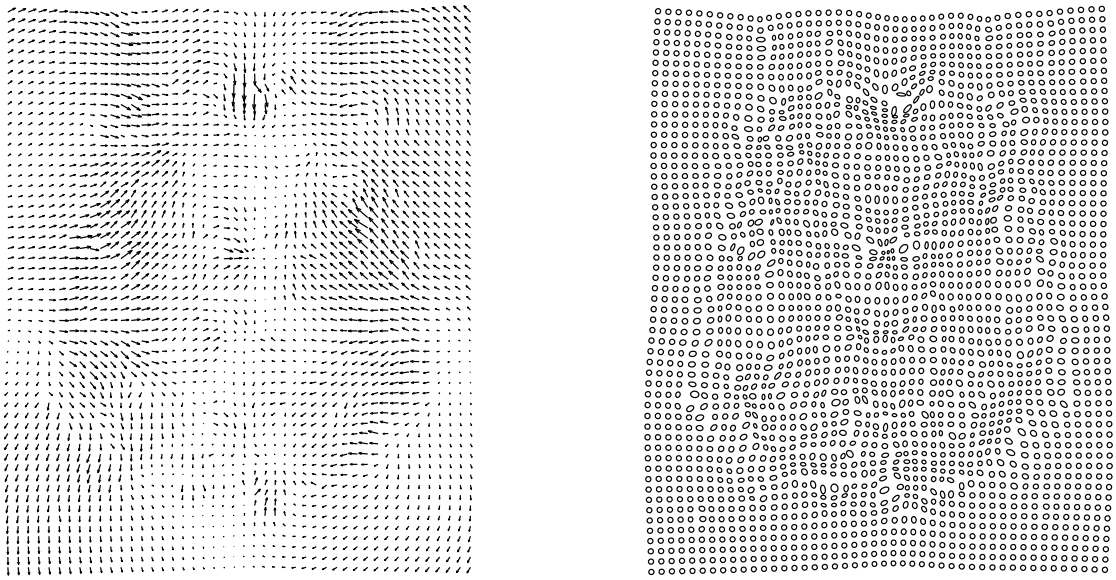

Figure 1.1: The term "deformation-based morphometry" will be used to describe methods of studying the positions of structures within the brain (left), whereas the term "tensor-based morphometry" will be used for methods that look at local shapes (right). Currently, the main application of tensor-based morphometry involves using the Jacobian determinants to examine the relative volumes of different structures. However, there are other features of the Jacobian matrices that could be used, such as those representing elongation and contraction in different directions. The arrows in the panel on the left show absolute displacements after making a global correction for rotations and translations, whereas the ellipses on the right show how the same circles would be distorted in different parts of the brain.

further characterised using a number of statistical procedures.

The terms deformation-based and tensor-based morphometry will be used to denote methods of studying brain shape that are based on deformation fields. When comparing groups, deformation-based morphometry (DBM) uses deformation fields to identify differences in the relative positions of structures within subjects' brains. Tensor-based morphometry (TBM) refers to those methods that identify differences in the local shape of brain structures (see Figure 1.1).

Characterisation using DBM can be global, pertaining to the entire field as a single observation, or can proceed on a voxel by voxel basis to make inferences about regionally specific positional differences. This simple approach to the analysis of deformation fields involves treating them as vector fields representing absolute displacements. However in this form, in addition to shape information, the vector fields also contain information on position and size that is likely to confound an analysis. Much of the confounding information is first removed by global rotations, translations and a zoom of the fields (Bookstein, 1997a).

DBM can be applied on a global scale to simply identify whether there are significant differences in overall shapes (based on a small number of parameters) among the brains of different populations. Generally, a single multi-variate test is performed using parameters describing the deformations - often after parameter reduction with singular value decomposition. The Hotelling's

T^2 statistic can be used for such simple comparisons between two groups of subjects (Bookstein, 1997a; Bookstein, 1999), but for more complex experimental designs, a multi-variate analysis of covariance can be used to identify differences via the Wilk's λ statistic.

An alternative approach to DBM involves producing a statistical parametric map that locates any regions of significant positional differences among the groups of subjects. An example of this approach involves using a voxel-wise Hotelling's T^2 test on the vector field describing the displacements at each and every voxel (Thompson & Toga, 1999; Gaser *et al.*, 1999). The significance of any observed differences can be assessed by modelling the statistic field as a T^2 random field (Cao & Worsley, 1999). Note that this approach does not directly localise brain regions with different shapes, but rather identifies those brain structures that are in relatively different positions.

If the objective is to localise structures whos shapes differ among groups, then some form of tensor-based morphometry is required to produce statistical parametric maps of regional shape differences. A deformation field that maps one image to another can be considered as a discrete vector field. By taking the gradients at each element of the field, a Jacobian matrix field is obtained, in which each element is a tensor describing the relative positions of neighbouring elements. Morphometric measures derived from such a tensor field can be used to locate regions with different shapes. The field obtained by taking the determinants at each point gives a map of structural volumes relative to those of a reference image (Freeborough & Fox, 1998; Gee & Bajcsy, 1999). Statistical parametric maps of these determinant fields can then be used to compare the anatomy of groups of subjects. A number of other measures derived from tensor fields have been used by other researchers, and these are described by Thompson and Toga (1999).

Another form of morphometry involves examining the local composition of brain images. Grey and white matter voxels can be identified by image segmentation, before applying morphometric methods to study the spatial distribution of the tissue classes. These techniques will be referred to as voxel-based morphometry (VBM). Currently, the difficulty of computing very high resolution deformation fields (required for TBM at small scales) makes voxel-based morphometry a simple and pragmatic approach to addressing small scale differences that is within the capabilities of most research units.

To summarise, computational neuroanatomic techniques can either use the deformation fields themselves or use these fields to normalise images that are then entered into an analysis of regionally specific differences. In this way, information about overall shape (deformations fields) and residual anatomic differences inherent in the data (registered images) can be partitioned.

1.2 Overview of Chapters

The remaining chapters of this thesis are organised as follows.

Rigid Body Registration

Rigid body registration is one of the simplest forms of image registration, so this chapter provides an ideal framework for introducing some of the concepts that will be used by the more complex registration methods described in later chapters. The shape of a human brain changes very little

with head movement, so rigid body transformations can be used to model different head positions of the same subject. Registration methods described in this chapter include within modality, or between different modalities such as PET and MRI. Matching of two images is performed by finding the rotations and translations that optimise some mutual function of the images. Within modality registration generally involves matching the images by minimising the sum of squared difference between them. For between modality registration, the matching criterion needs to be more complex. A method for co-registering brain images of the same subject that have been acquired in different modalities is presented. The basic idea is that instead of matching two images directly, one performs intermediate within modality registrations to two template images that are already in register. One can use a least squares minimisation to determine the affine transformations that map between the templates and the images. By incorporating suitable constraints, a rigid body transformation that directly maps between the images can be extracted from these more general affine transformations. A further refinement capitalises on the implicit normalisation of both images into a standard space. This facilitates partitioning both original images into homologous tissue classes. Once extracted, the partitions are jointly matched further increasing the accuracy of the co-registration.

Image Warping with Basis Functions

This chapter describes the steps involved in registering images of different subjects into roughly the same co-ordinate system, where the co-ordinate system is defined by a template image (or series of images). The method only uses up to a few hundred parameters, so can only model global brain shape. It works by estimating the optimum coefficients for a set of bases, by minimising the sum of squared differences between the template and source image, while simultaneously maximising the smoothness of the transformation using a *maximum a posteriori* (MAP) approach. In order to adopt the MAP approach, it is necessary to have estimates of the likelihood of obtaining the fit given the data, which requires prior knowledge of spatial variability, and also knowledge of the variance associated with each observation. True Bayesian approaches assume that the variance associated with each voxel is already known, whereas the approach developed here is a type of Empirical Bayesian method, which attempts to estimate this variance from the residual errors. Because the registration is based on smooth images, correlations between neighbouring voxels are considered when estimating the variance. This makes the same approach suitable for the spatial normalisation of both high quality MR images, and low resolution noisy PET images. A fast algorithm has been developed that utilises Taylor's Theorem and the separable nature of the basis functions, meaning that most of the nonlinear spatial variability between images can be automatically corrected within a few minutes.

The approach begins by matching the images using an affine transformation. Unlike Chapter 2 – where the images to be matched together are from the same subject – zooms and shears are needed to register heads of different shapes and sizes. Knowledge of the variability of head sizes is included within a Bayesian framework in order to increase the robustness and accuracy of the method. Following this step, gross differences in head shapes, that can not be accounted for by affine normalisation alone, are corrected by a nonlinear spatial normalisation procedure. In order to reduce the number of parameters to be estimated, the nonlinear warps are described by a linear combination of low spatial frequency discrete cosine transform basis functions. Regularisation of the problem involves biasing the warps to be smooth by simultaneously minimising their

membrane energy.

High Dimensional Image Warping

This chapter is also about warping brain images of different subjects to the same stereotactic space. However, unlike Chapter 3, this method uses thousands or millions of parameters, so is potentially able to obtain much more precision. A high dimensional model is used, whereby a finite element approach is employed to estimate translations at the location of each voxel in the template image. Bayesian statistics are used to obtain a *maximum a posteriori* (MAP) estimate of the deformation field. The validity of any registration method is largely based upon the prior knowledge about the variability of the estimated parameters. In this approach it is assumed that the priors should have some form of symmetry, in that priors describing the probability distribution of the deformations should be identical to those for the inverses (i.e., warping brain A to brain B should not be different probablistically from warping B to A). The fundamental assumption is that the probability of stretching a voxel by a factor of n is considered to be the same as the probability of shrinking n voxels by a factor of n^{-1}. The penalty function of choice is based upon the singular values of the Jacobian matrices having log-normal distributions, which enforces a continuous one-to-one mapping. A gradient descent algorithm is presented that incorporates the above priors in order to obtain a MAP estimate of the deformations. Further consistency is achieved by registering images to their "averages", where this average is one of both intensity and shape.

Segmentation

A tissue classification method was originally developed to be part of the between modality registration procedure described in Chapter 2, but the classification results are also useful for various types of morphometry, as well as having potential applications in other registration techniques. This chapter describes a method of segmenting MR images into different tissue classes, using a modified Gaussian Mixture Model. By knowing the prior spatial probability of each voxel being grey matter, white matter or cerebro-spinal fluid, it is possible to obtain a more robust classification. In addition, a step for correcting intensity non-uniformity is also included, which makes the method more applicable to images corrupted by smooth intensity variations. Evaluations of the method show that the non-uniformity correction improves the segmentation of images containing this artifact.

Morphometry

The chapter on morphometry covers three principle morphometric methods, that will be called *voxel-based*, *deformation-based* and *tensor-based* morphometry.

At its simplest, voxel-based morphometry (VBM) involves a voxel-wise comparison of the local concentration of grey matter between two groups of subjects. The procedure is relatively straightforward, and involves spatially normalising high resolution MR images from all the subjects in the study into the same stereotactic space. This is followed by segmenting the grey matter from the spatially normalised images, and smoothing these grey-matter segments. Voxel-wise parametric

statistical tests are performed which compare the smoothed grey-matter images from the groups. Corrections for multiple comparisons are made using the theory of Gaussian random fields. This chapter describes the steps involved in VBM, and provides evaluations of the assumptions made about the statistical distribution of the data.

Deformation-based morphometry (DBM) is a method for identifying macroscopic anatomical differences among the brains of different populations of subjects. The method involves spatially normalising the structural MR images of a number of subjects so that they all conform to the same stereotactic space. Multivariate statistics are then applied to the parameters describing the estimated nonlinear deformations that ensue. To illustrate the method, the gross morphometry of male and female subjects are compared. Brain asymmetry, the effect of handedness, and the interactions among these effects are also assessed.

Tensor-based morphometry (TBM) is introduced as a method of identifying regional structural differences from the gradients of deformations fields. Deformation fields encode the relative positions of different brain structures, but local shapes (such as volumes, lengths and areas) are encoded in their gradients (Jacobian matrix field). Various functions of these tensor-fields can be used to characterise shape differences.