Chapter 6

Morphometry

John Ashburner & Karl J. Friston
The Wellcome Dept. of Imaging Neuroscience,
12 Queen Square, London WC1N 3BG, UK.

Contents

Abstract

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

Deformation-based morphometry (DBM) is a method for identifying macroscopic anatomical differences among the brains of different populations of subjects. The method involves spatially normalizing the structural MR images of a number of subjects so that they all conform to the same stereotaxic space. Multivariate statistics are then applied to the parameters describing the estimated nonlinear deformations that ensue.

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 characterize shape differences.

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 straight-forward, and involves spatially normalizing high-resolution MR images from all the subjects in the study into the same stereotaxic space. This is followed by segmenting the grey matter from the spatially normalized 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.

6.1 Introduction

The morphometric methods described in this chapter relate to ways of statistically identifying and characterizing structural differences among populations, or for finding correlations between brain shape and, for example, disease severity. A large number of approaches for characterizing differences in the shape and neuro-anatomical 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.

Studies of brain shape have been carried out by many researchers on a number of different populations, including patients with schizophrenia, autism, dyslexia and Turner's syndrome. Often, the morphometric measurements used in these studies have been obtained from brain regions that can be clearly defined, resulting in a wealth of findings pertaining to these particular measurements. The measures are typically volumes of unambiguous structures such as the hippocampi or the ventricles. However, there are a number of morphometric features that may be more difficult to quantify by inspection, meaning that many observable structural differences may be overlooked. This chapter describes morphometric approaches that are not biased to one particular structure or tissue and give even-handed and comprehensive assessments of anatomical differences throughout the brain.

The approaches described here require the images of multiple subjects to be registered together by some form of spatial normalization. The primary result of spatially normalizing a series of images is that they all conform to the same stereotaxic 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 further characterized 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 6.1).

Characterization 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

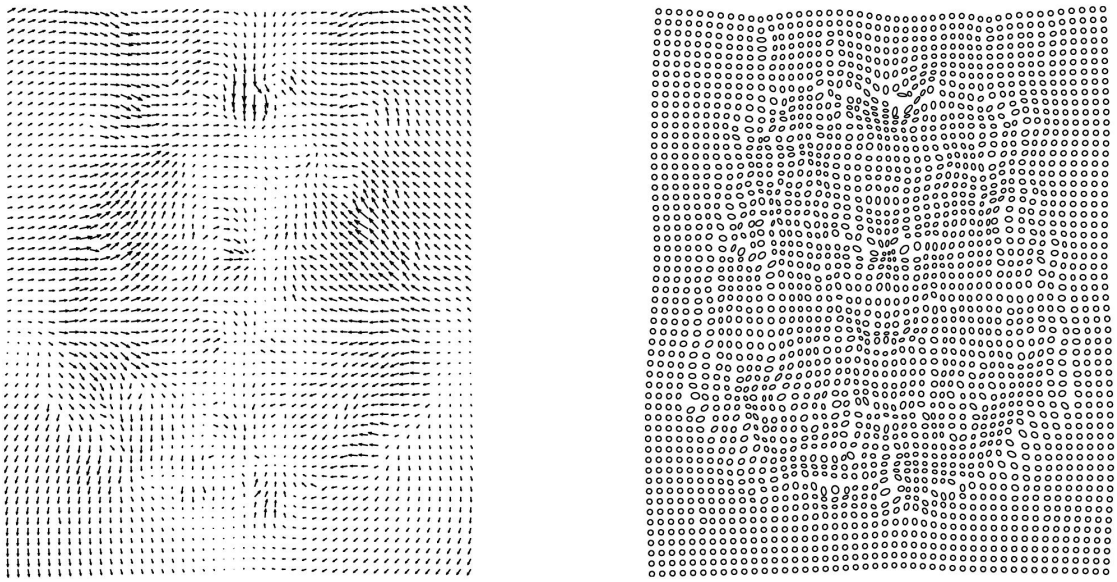

Figure 6.1: The term "deformation-based morphometry" will be used to describe methods of studying the positions of structures within the brain (left), whereas the term "tensor-based morphometry" will be used for methods that look at local shapes (right). Currently, the main application of tensor-based 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.

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 [4].

When applied on a global scale, DBM simply identifies 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 [4, 6, 13], 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 [23, 17]. The significance of any observed differences can be assessed by modeling the statistic field as a T^2 random field [9]. Note that this approach does not directly localize brain regions with different shapes, but rather identifies those brain structures that are in relatively different positions. The locations of the results will be highly dependent upon how the positions and sizes of the brains are standardized.

If the objective is to localize structures whose shapes differ among groups, then some form of tensor-based morphometry is preferable for producing 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 neighboring 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 [14, 18, 11, 10]. 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 also been used by other researchers [23, 22].

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.

Morphometric methods have a number of different aims. They can be used for localizing significant structural differences among populations, or for showing that overall brain structure is related to some effect of interest. When testing the overall brain structure, multivariate statistical methods are used to analyze groups of parameters for the whole brain (e.g., the deformation-based morphometry described in this chapter). The result of the forms of morphometry that localize structural differences would typically be a statistical parametric map [16] of regional differences. Statistical parametric maps (SPMs) can be derived from uni-variate data where there is a single variable at each voxel (e.g., the voxel-based morphometry method described here), or from multi-variate data, where there are several different variables at each voxel (e.g., tensor-based morphometry).

Another use for morphometric methods is for characterizing essential differences, or for producing some form of classification. Linear methods such as canonical correlation analysis, or nonlinear classification methods can be used for these purposes. This chapter will be restricted

to simple linear classification methods, which model data as multivariate normal distributions. Nonlinear classification methods can assume more complex distributions for the data, but they tend to be much more computationally expensive.

6.2 Deformation-Based Morphometry

Deformation-based morphometry (DBM) is a characterization of the differences in the vector fields that describe global or gross differences in brain shape. These vector fields are the deformation fields used to effect nonlinear variants of spatial normalization, when one of the images is a template that conforms to some standard anatomical space.

Spatially normalizing a series of image volumes to match the same template will result in a series of deformation fields – one for each image. A deformation field can be considered as a continuous 3D vector field, containing a three element vector at each point. After spatial normalization, what each deformation field contains is a mapping from each voxel in the template image[1] to its corresponding voxel in the un-warped image. A simple representation of a deformation field is a series of three volumes of the same dimensions as the template image. The lattice location in one of these volumes corresponds to the same location in the template. The contents of a lattice location in the three volumes constitutes a vector that points to the co-ordinate of the equivalent structure in the un-normalized image. Therefore, the contents of voxel $\mathbf{x}$ in each of the deformation fields will be the co-ordinate of the same structure in each of the un-warped images. This allows statistical analysis methods to be used for comparing the positions of structures in different subject's images. The shape of an object is defined by the relative positions of its components, so in studying the relative positions of structures within a brain, we are actually finding out about the brain's shape.

In addition to encoding the shape of a subject's brain, a deformation field also encodes its pose and size. To make inferences about shape alone, these confounding components need to be removed.

Because deformation fields are multivariate, standard multivariate statistical techniques can be employed to estimate the nature of the differences and to make inferences about them. The endpoint of this form of DBM is a p value, derived via the Wilk's λ or Hotelling's T^2 statistic, pertaining to the significance of the effect and one or more canonical vectors, or deformations, that characterize their nature. These results can be obtained using multivariate analysis of covariance (MANCOVA) and canonical correlation analysis (CCA) respectively. Before proceeding with the multivariate analyses, it is necessary for the information in the deformations to be compacted to just a few parameters. This is usually done by generating "eigen-warps" (principal modes of variation of the warps [18]) using singular value decomposition.

6.2.1 Extracting Shape Information

Given a series of deformation fields, it is necessary to decompose each deformation into components relating to global position, orientation and size[2] (uninteresting components), and shape (the components of interest). Each field provides a mapping from points in the template to points in the source image, allowing well known landmark methods [5, 13] to be used to extract the size and positional information. The extracted measures are such that the remaining shapes minimize the

[1] Actually, it is a mapping from each voxel in a spatially normalized image to its voxel in the original un-warped image. We are assuming that the spatial normalization is 100% accurate, such that each voxel in the template exactly matches the same voxel in the normalized image.

[2] Although the total size of a brain is also often of interest.

squared Procrustes distance between the template and images. Rather than basing the procedure on a few landmarks, all the elements of the deformation field corresponding to voxels within the brain are considered. This involves first determining translations by computing centers of mass:

$$\bar{\mathbf{x}} = \frac{\sum_{i=1}^{I} \mathbf{x}_i w_i}{\sum_{i=1}^{I} w_i}$$
$$\bar{\mathbf{y}} = \frac{\sum_{i=1}^{I} \mathbf{y}_i w_i}{\sum_{i=1}^{I} w_i}$$

where $\mathbf{x}_i$ is the co-ordinate of the ith voxel of the template, $\mathbf{y}_i$ is the location that it maps to, and w_i is a weighting for that element. Non-brain voxels are typically given zero weighting. The rotations are computed from the cross-covariance matrix ($\mathbf{C}$) between the elements and deformed elements (after removing the effects of position):

$$c_{jk} \propto \sum_{i=1}^{I} w_i (x_{ij} - \bar{x}_j)(y_{ik} - \bar{y}_k)$$

The 3×3 matrix $\mathbf{C}$ is decomposed using singular value decomposition to give three matrices, $\mathbf{U}$, $\mathbf{S}$ and $\mathbf{V}$ (such that $\mathbf{C} = \mathbf{USV}^T$, where $\mathbf{U}$ and $\mathbf{V}$ are unitary, and $\mathbf{S}$ is a diagonal matrix). The rotation matrix ($\mathbf{R}$) can then be reconstituted from these matrices by $\mathbf{R} = \mathbf{UV}^T$. If size effects are to be removed, then finally, moments around the centers are used to correct for relative size differences (z):

$$z = \sqrt{\frac{\sum_{j=1}^{3} \sum_{i=1}^{I} (x_{ij} - \bar{x}_j)^2 w_i}{\sum_{j=1}^{3} \sum_{i=1}^{I} (y_{ij} - \bar{y}_j)^2 w_i}}$$

6.2.2 Multivariate Analysis of Covariance

Following this, a data matrix is generated, where each row contains a corrected deformation field for one subject represented by a single vector. For multivariate analysis, it is necessary to have a representation of each deformation that is parameterized by far fewer variables than the number of subjects included in the study. Singular value decomposition can be used to compact this information, such that most of the variance of the nonlinear deformations can be represented by only a few parameters for each subject. MANCOVA is then used to make inferences about the effects of interest (i.e., provide p values).

A multivariate analysis of covariance (MANCOVA) assumes that there are several dependent variables for each observation, where an observation refers to a collection of data for a subject. This data can be represented by an $M \times I$ matrix $\mathbf{X}$, where M is the number of subjects in the analysis, and I is the number of variables for each subject.

The columns of $\mathbf{X}$ are modeled by a linear combination of basis functions. Some of these basis functions represent effects that are not considered interesting in the study, but may still be significant. For example, linear age effects may confound a study of handedness. If a left-handed group is not perfectly age matched with a right-handed group, then findings that are actually due to age could be attributed to handedness differences. In addition, the inclusion of confounding effects (such as age) in a model can result in a better fit to the data, possibly making the test more sensitive to the effects of interest.

Confounding effects are modeled by an $M \times K$ design matrix $\mathbf{G}$. Each column of $\mathbf{G}$ can be a vector of covariates (e.g., the age of each subject), or alternatively can be arranged in blocks (e.g., there may be a column containing ones for left-handed subjects, and zeros for right-handed). In

almost all cases, a column of ones is included in order to model a mean effect. First of all, any variance in the data that could be attributed to the confounds is removed by an orthogonalization step:

$$\mathbf{X_a} = \mathbf{X} - \mathbf{G}\left(\mathbf{G}^T\mathbf{G}\right)^{-1}\mathbf{G}^T\mathbf{X} \tag{6.1}$$

Similarly, the effects of interest are modeled by an $M \times J$ design matrix $\mathbf{Y}$. The columns of $\mathbf{Y}$ can represent group memberships, where elements contain a one for subjects who are in a particular group, or a zero if they do not. Alternatively, the columns may contain covariates of interest such as disease severity for each subject. The columns in this design matrix are orthogonalized with respect to matrix $\mathbf{G}$:

$$\mathbf{Y_a} = \mathbf{Y} - \mathbf{G}\left(\mathbf{G}^T\mathbf{G}\right)^{-1}\mathbf{G}^T\mathbf{Y} \tag{6.2}$$

A MANCOVA involves assessing how the predictability of the observations change when the effects of interest are discounted. This is based on the distributions of the residuals, which are assumed to be multi-normal. The statistic is related to the determinants of the covariance matrices describing these distributions. In practice, the residual sum of squares and products (SSP) matrix ($\mathbf{W}$), is compared to the SSP matrix of the fitted effects ($\mathbf{B}$). These matrices are obtained by:

$$\begin{aligned}\mathbf{T} &= \mathbf{Y_a}({\mathbf{Y_a}}^T\mathbf{Y_a})^{-1}{\mathbf{Y_a}}^T\mathbf{X_a}\\ \mathbf{B} &= \mathbf{T}^T\mathbf{T}\\ \mathbf{W} &= (\mathbf{X_a}-\mathbf{T})^T(\mathbf{X_a}-\mathbf{T})\end{aligned}$$

The statistic is called *Wilk's Lambda* (Λ), and is based upon the ratios of the determinants:

$$\Lambda = \frac{|\mathbf{W}|}{|\mathbf{B}+\mathbf{W}|}$$

The Wilk's Lambda statistic can range between zero and one, where a value of one suggests no relationship between the effects of interest and the data, and a value of zero indicates a perfect relationship. This statistic is transformed to a χ^2 statistic (with IJ degrees of freedom under the null hypothesis) using the approximation of Bartlett [3]:

$$\chi^2 \approx -(M-K-(I+J+1)/2)log_e(\Lambda)$$

Finally, the cumulative χ^2 distribution function is used to make inferences about whether the null hypothesis (that there is no difference between the distributions) can be rejected.

This multivariate approach fails when the number of variables approaches the number of subjects ($M-K-(I+J+1)/2$ approaches zero, or becomes negative). In many situations, it is necessary to regularize the problem by reducing the number of variables with respect to the number of subjects. One way of doing this involves using singular value decomposition to decompose the original data matrix $\mathbf{X}$ into unitary matrices $\mathbf{U}$ and $\mathbf{V}$, and diagonal matrix $\mathbf{S}$, such that $\mathbf{X} = \mathbf{USV}^T$. The diagonal elements of $\mathbf{S}$ are arranged in decreasing importance, so it is possible to reconstruct an approximation of $\mathbf{X}$ using only the first L diagonal elements of $\mathbf{S}$ and the first L columns of $\mathbf{U}$ and $\mathbf{V}$, such that $\mathbf{X} \simeq \mathbf{U}^*\mathbf{S}^*{\mathbf{V}^*}^T$. The MANCOVA would be performed using $\mathbf{U}^*$ (the first L columns of $\mathbf{U}$).

[3] This is true provided the matrices do not contain any redundant columns, otherwise matrix pseudo-inverses are required in the computations, and the numbers of columns replaced by the matrix ranks.

6.2.3 Canonical Correlation Analysis

To characterize differences among groups of subjects, one usually uses CCA, which is a device that finds the linear combination of dependent variables (in this case the deformations) that is maximally correlated with the explanatory variables (e.g., male vs. female). In a simple case of one categorical explanatory variable (e.g., sex) this will be the deformation field that best discriminates between the groups. Note that this is not the same as simply subtracting the deformation fields of two groups. This is because (i) the MANCOVA includes the effects of confounds that are removed and (ii) some aspects of the deformations may be less reliable than others (CCA gives deformations that explicitly discount error in relation to predicted differences). The canonical deformations can either be displayed directly as deformation fields, or can be applied to some image to "caricature" the effect detected.

Canonical correlation analysis (CCA) is used to measure the strength of association between two sets of variables. In this case, the variables are $\mathbf{X_a}$ and $\mathbf{Y_a}$, which are the data and design matrix from Section 6.2.2 after having been orthogonalized with respect to a set of confounding effects. The first canonical variate pair is the linear combination of columns of $\mathbf{X_a}$ and the linear combination of columns of $\mathbf{Y_a}$, that has the maximum correlation. The second canonical variate pair consists of linear combinations that maximize the correlation subject to the constraint that they are orthogonal to the first pair of canonical variables. Similarly, all subsequent pairs maximize the correlations and are orthogonal to all the previous pairs.

The weights used to determine the linear combinations are derived from the unitary matrices ($\mathbf{U}$ and $\mathbf{V}$) obtained by singular value decomposition:

$$\mathbf{USV}^T = \left(\mathbf{X_a}^T\mathbf{X_a}\right)^{-\frac{1}{2}}\left(\mathbf{X_a}^T\mathbf{Y_a}\right)\left(\mathbf{Y_a}^T\mathbf{Y_a}\right)^{-\frac{1}{2}}$$

Then the weights ($\mathbf{A}$ and $\mathbf{B}$) are derived by:

$$\mathbf{A} = \left(\mathbf{X_a}^T\mathbf{X_a}\right)^{-\frac{1}{2}}\mathbf{U} \text{ and } \mathbf{B} = \left(\mathbf{Y_a}^T\mathbf{Y_a}\right)^{-\frac{1}{2}}\mathbf{V}$$

The canonical variate pairs are obtained from the columns of $\mathbf{X_aA}$ and $\mathbf{Y_aB}$.

When the number of variables approaches the number of observations, the problem needs to be regularized. This can be done by computing the canonical variates from data that has been compacted using singular value decomposition as in Section 6.2.2. If $\mathbf{X}$ has been decomposed such that it can be approximated by $\mathbf{X} \simeq \mathbf{U}^*\mathbf{S}^*{\mathbf{V}^*}^T$, and canonical correlation analysis performed on $\mathbf{U}^*$ and $\mathbf{Y}$ to give weight matrices $\mathbf{A}^*$ and $\mathbf{B}^*$, then the weights to be applied to the original data ($\mathbf{X}$) can be reconstructed by $\mathbf{V}^*{\mathbf{S}^*}^{-1}\mathbf{A}^*$.

Figure 6.2 shows CCA as it would be used to graphically describe the differential features of three groups. However, it can also be used to aid classification. Once derived, the same weighting matrix ($\mathbf{A}$) can be applied to new data-sets that were not involved in its derivation. If there are only two groups involved in a study, then CCA can be used directly to assign group memberships to new observations[4]. With more groups, CCA serves as a graphical aid for assigning new observations to the groups.

6.3 Tensor-Based Morphometry

There is a near infinite number of ways in which the shapes of brains can differ among populations of subjects. Many thousands or millions of parameters are required to precisely describe the shape

[4] Providing that the signs of the canonical variates are adjusted accordingly.

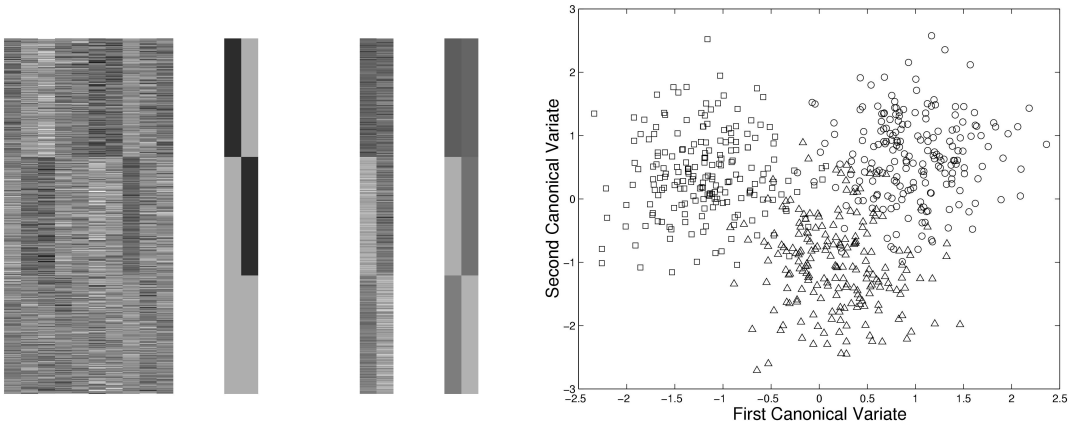

Figure 6.2: This figure illustrates canonical correlation analysis using simulated data. The two matrices on the left are a 600×10 data matrix ($\mathbf{X_a}$) and a 600×2 design matrix ($\mathbf{Y_a}$) after centering by orthogonalizing with respect to a column of ones as described in Equations 6.1 and 6.2. The design matrix represents a partitioning of the data into three groups. The canonical variates for $\mathbf{X_a}$ and $\mathbf{Y_a}$ are shown in the third and fourth matrices ($\mathbf{X_aA}$ and $\mathbf{Y_aB}$). The left-hand columns of these two matrices contain the first canonical variate pair, whereas the second pair are in the right-hand columns. The graph on the right shows the two columns of $\mathbf{X_aA}$ plotted against each other, where the different symbols used represent memberships of the three different groups.

of a brain at the resolution of a typical structural MR image. Given, say, 10 schizophrenic and 10 control brain images, there are lots of ways of inventing a measure that would differentiate between the groups. In most cases though, this measure will not provide any distinguishing power in a comparison between further groups of schizophrenics and controls. In other words, the measure would be specific to the subjects included in the study, and not generalizable to the populations as a whole. It is therefore not feasible to use methods that try to detect *any* difference. One has to be specific about the types of differences that are searched for.

The multivariate methods described in Section 6.2 require that the deformations are parameterized by only a few "eigen-warps", on which the statistical tests are performed. This only allows inferences about those eigen-warps included in the analyses, so much of the information in the deformation fields is lost.

A more useful analysis involves identifying focal differences in the form of a statistical parametric map (SPM). SPMs of univariate statistical measures often allow relatively simple questions to be addressed, such as where is there significantly more of a particular measure that happens to correlate with a particular effect of interest. Standard parametric statistical procedures (t-tests and F-tests) can be used to test the hypotheses within the framework of the general linear model (GLM – see Chapter 7), whereby a vector of observations is modeled by a linear combination of user specified regressors [16]. The GLM is a flexible framework that allows many different tests to be applied, ranging from group comparisons and identifying differences that are related to specified covariates such as disease severity or age, to complex interactions between different effects of interest.

Performing comparisons at each voxel results in many statistical tests being done. Without any correction, the number of false-positive results would be proportional to the number of independent tests. A Bonferroni correction would be applied if the tests were independent, but this is not normally the case because of the inherent spatial smoothness of the data. In practice,

the effective number of independent statistical tests is determined using Gaussian Random Field (GRF) theory [15, 25] (see Chapters 14 and 15). By using GRF theory, a correction for multiple dependent comparisons can be made to produce the appropriate rate of false-positive results.

SPMs can also be obtained from the results of voxel-wise multi-variate tests. Instead of one variable per voxel of a subject, multivariate tests could effectively involve two or more variables. Following the voxel-wise multivariate tests, similar corrections based on GRF theory can be applied as in the univariate case. GRF theory has not yet been extended for Wilk's Λ fields, so approximations would need to be made that involve transforming the resulting Wilk's Λ fields to random fields of other statistics, such as χ^2 or F fields. Subsequent processing would then need to assume that the transformed fields have the same properties as true χ^2 or F fields. Another useful multivariate measure is the Hotelling's T^2 statistic, for which GRF theory has recently been extended [9].

The objective of TBM is to localize regions of shape differences among groups of brains, based on deformation fields that map points in a template (x_1, x_2, x_3) to equivalent points in individual source images (y_1, y_2, y_3). In principle, the Jacobian matrices of the deformations (a 2nd order tensor field relating to the spatial derivatives of the transformation) should be more reliable indicators of local brain shape than absolute deformations. Absolute deformations represent positions of brain structures, rather than local shape, and need to be quantified relative to some arbitrary reference position.

A Jacobian matrix contains information about the local stretching, shearing and rotation involved in the deformation, and is defined at each point by:

$$\mathbf{J} = \begin{bmatrix} \partial y_1/\partial x_1 & \partial y_1/\partial x_2 & \partial y_1/\partial x_3 \\ \partial y_2/\partial x_1 & \partial y_2/\partial x_2 & \partial y_2/\partial x_3 \\ \partial y_3/\partial x_1 & \partial y_3/\partial x_2 & \partial y_3/\partial x_3 \end{bmatrix}$$

A simple form of TBM involves comparing relative volumes of different brain structures, where the volumes are derived from Jacobian determinants at each point (see Figures 6.3 and 6.4). Simple uni-variate statistics (t- or F- tests) can then be used to make inferences about regional volume differences among populations. This type of morphometry is useful for studies that have specific questions about whether growth or volume loss has occurred.

When many subjects are included in a study, a potentially more powerful form of TBM can be attained using multi-variate statistics on other measures derived from the Jacobian matrices. This use of multivariate statistics does not only test for volumetric differences, but indicates whether there are any differences among lengths, areas and the amount of shear. It may therefore be useful when there is no clear hypothesis about the nature of the differences, as may be the case when studying the effects of maturation on the human brain. This form of morphometry should be able to identify shape differences even when volumes are the same.

Because a Jacobian matrix encodes both local shape (zooms and shears) and orientation, it is necessary to remove the latter before making inferences. According to the polar decomposition theorem [21], a non-singular Jacobian matrix can be decomposed into a rotation matrix ($\mathbf{R}$) and a symmetric positive definite matrix ($\mathbf{U}$ or $\mathbf{V}$), such that $\mathbf{J} = \mathbf{RU} = \mathbf{VR}$. Matrices $\mathbf{U}$ and $\mathbf{V}$ (called the *right* and *left stretch* tensors respectively) are derived by $\mathbf{U} = \left(\mathbf{J}^T\mathbf{J}\right)^{1/2}$ and $\mathbf{V} = \left(\mathbf{J}\mathbf{J}^T\right)^{1/2}$. Matrix $\mathbf{R}$ is then given by $\mathbf{R} = \mathbf{J}\mathbf{U}^{-1}$ or $\mathbf{R} = \mathbf{V}^{-1}\mathbf{J}$. For a purely rigid body transformation, $\mathbf{U} = \mathbf{V} = \mathbf{I}$ (the identity matrix). Similarly, if $\mathbf{R} = \mathbf{I}$, then the deformation can be considered as *pure strain*. Deviations of $\mathbf{U}$ or $\mathbf{V}$ away from $\mathbf{I}$ indicate a shape change, which can be represented by a strain tensor $\mathbf{E}$. Depending upon the reference co-ordinate system, the strain tensor is referred to as either a *Lagrangean* or an *Eulerian* strain tensor. When the strain tensor is derived from $\mathbf{U}$, it is referred to as a Lagrangean strain tensor, whereas when it is derived from $\mathbf{V}$, it is Eulerian. Spatial normalization of a series of source images involves

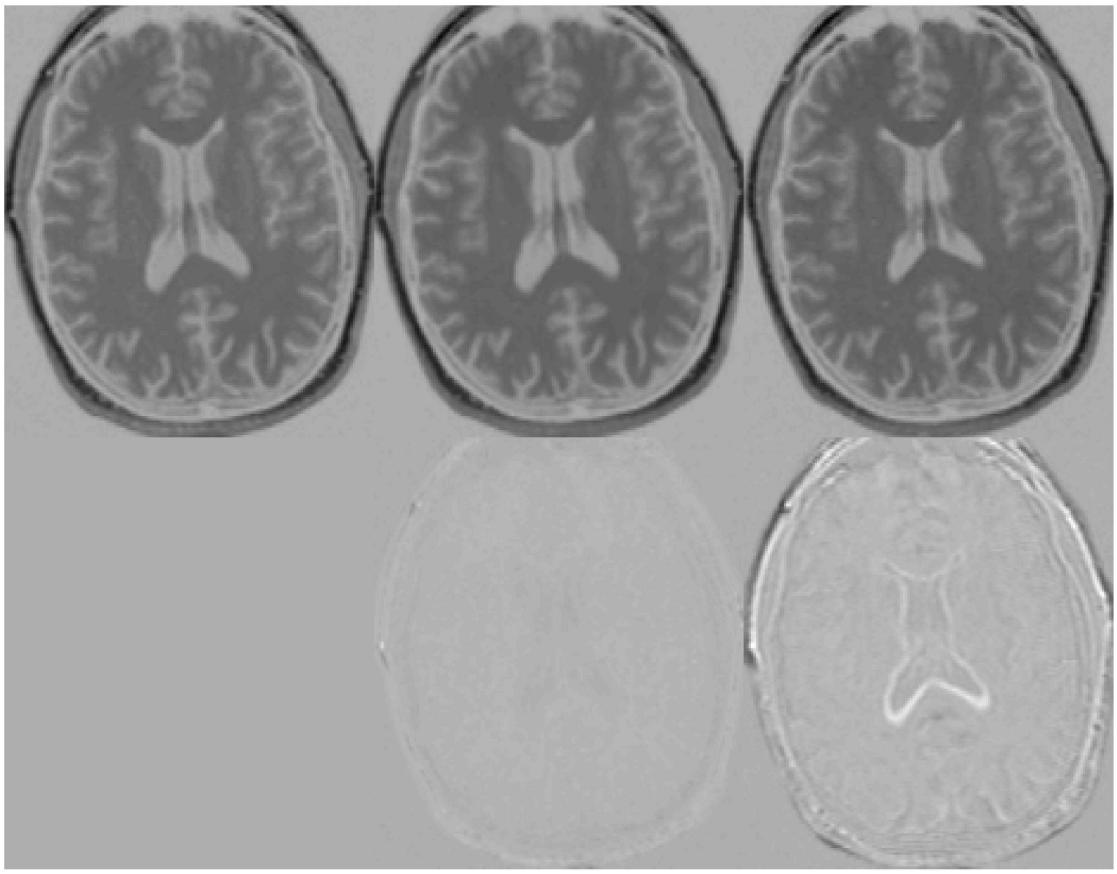

Figure 6.3: This figure illustrates warping together a pair of serial scans of the same Alzheimers subject, using the three dimensional method presented in Chapter 4. The data were provided by the Dementia Research Group, National Hospital for Neurology and Neurosurgery, London, UK. From left to right, the top row shows a slice from the late scan (used as the template), the early scan after rigid body registration followed by warping, and the early scan after rigid registration only. The bottom row shows the difference between the late and early scan, with and without warping. All images are displayed with the same intensity scaling. With only rigid registration, the difference between the early and late scan exhibits a difference at the edges of the ventricles (among other regions). This difference is greatly reduced by including warping in the spatial transformation. The information contained in the difference image is transferred to the deformation fields. One informative feature of the deformation fields are the Jacobian determinants, which reflect volume changes. These are shown in Figure 6.4.

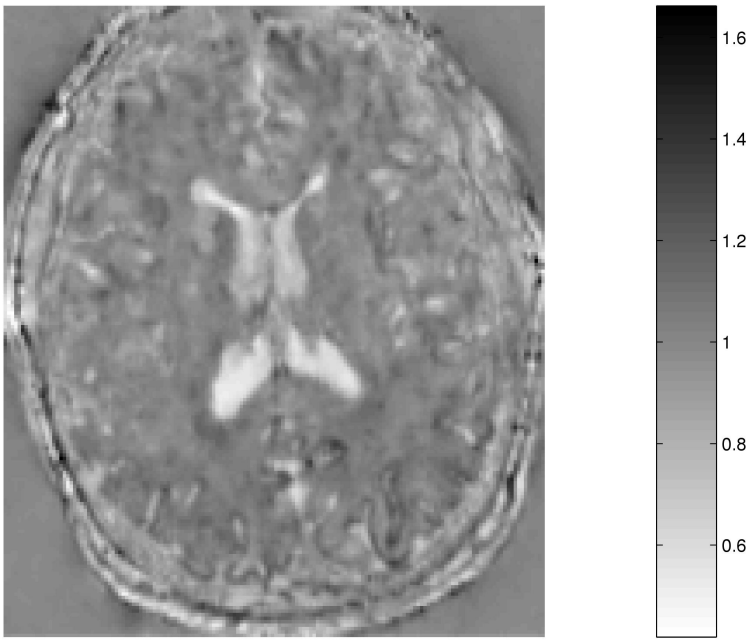

Figure 6.4: This figure illustrates the volume changes estimated by warping together the images shown in Figure 6.3. The relative volumes are the Jacobian determinants of the deformation field. Smaller determinants are obtained when a region of the template maps to a smaller region in the source image. In this example, they represent regions that have expanded between the early and late scans. Regions where there are no measurable volume changes have Jacobian determinants with a value of one.

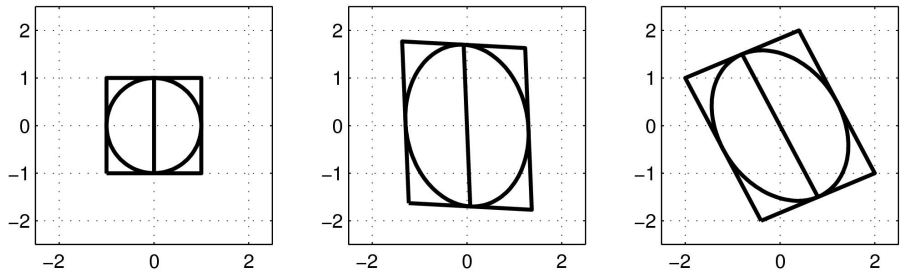

Figure 6.5: This figure illustrates polar decomposition, whereby the regular shape in the left-hand panel is transformed via shears and zooms to give the shape in the center, before being rotated to give the right-hand shape.

determining a mapping from each point in the template image, to corresponding points in the source images. To compare image shapes, it is necessary to derive measures of shape within the co-ordinate system of the template image, rather than within the different co-ordinate systems of the individual source images. Therefore, the Lagrangean framework should be used.

For any deformation, there is a whole continuum of ways of defining strain tensors, based on a parameter m. When m is non-zero, the family of Lagrangean strain tensors $\mathbf{E}$ are given by $\mathbf{E}^{(m)} = m^{-1}(\mathbf{U}^m - \mathbf{I})$. For the special case when m is zero, the strain tensor is given by $\mathbf{E}^{(0)} = ln(\mathbf{U})$, where the ln refers to a matrix logarithm. When m assumes values of -2, 0, 1 or 2, then the tensors are called the *Almansi*, *Hencky*, *Biot* or *Green* strain tensors respectively. For deformations derived using the methods in Chapter 4[5], the Hencky tensor may be the most appropriate measure of local shape, providing the priors exert enough influence on the resulting deformation fields. One advantage of basing measures on the Hencky tensor is that it allows relative changes in length, volume and area to be modeled using similar log-normal distributions. This would not be possible if the measures were modeled by normal distributions, because if a variable l is normally distributed, then l^3 is not. Also, it would not make sense to assume distributions that allow negative lengths or volumes, if these are explicitly prohibited by the warping algorithm.

In three dimensions, a strain tensor has six unique elements (as it is a symmetric 3×3 matrix). Within a TBM framework, multivariate statistics would be applied to these elements in order to localize volume, area and length differences, whereas DBM would involve performing a single, global multivariate test to compare the shapes. In principle, the difference between these approaches need not be quite so distinct. Rather than basing the tests either on information at each voxel, or on information from the whole brain, it is easy to see that forms of morphometry could be devised that are based on regional analyses. For example, DBM could be done region by region (after factoring out positional information), or multivariate TBM could be applied such that the tests include information from strain tensors in regions of neighboring lattice locations. Both of these approaches should produce similar results. The regions could be based on pre-defined structures on the template image. Alternatively, they could involve some weighted combination of information from the neighborhood of each voxel, with weights based on a function such as a Gaussian kernel, or some other windowing function.

[5] It should be noted that the sum of squares of the Hencky tensor elements is equivalent to the sum of squares of the logs of the singular values of $\mathbf{J}$, meaning that the priors described in Chapter 4 are effectively minimizing the squares of the Hencky tensor.

6.4 Voxel-Based Morphometry

A number of studies have already demonstrated structural brain differences among different patient populations using the technique of Voxel-Based Morphometry (VBM) [27, 26, 2], or a related method that involves parceling the images into different regions [19, 11].

VBM is another technique for producing SPMs of volumetric differences. Structural MR images of brains can differ among subjects in many ways. A useful measure of structural difference among populations is derived from a comparison of the local composition of different brain tissue types (e.g., grey matter, white matter etc). Voxel-based Morphometry (VBM) has been designed to be sensitive to these differences, while discounting positional and other large scale volumetric differences in gross anatomy. The technique involves spatially normalizing all the subjects' MR images to the same stereotaxic space, extracting the grey matter from the normalized images, smoothing, and finally performing a statistical analysis to localize, and make inferences about, group differences. The output from the method is an SPM showing regions where grey matter concentration differs significantly among the groups.

VBM was originally devised to detect cortical thinning in a way that was not confounded by volume changes of the sort that are characterized by classical volumetric analyses of large brain structures (e.g., the temporal lobe). It does this by removing positional and volume difference (down to a specified spatial scale) through spatial normalization. Differences in grey matter density are then detected by comparing the local intensities of grey matter maps after smoothing. Since its inception [27] VBM has become an established tool in morphometry being used to detect cortical atrophy and differences in slender white matter tracts. The classical perspective on VBM partitions volume changes into two spatial scales (i) macroscopic volume or shape differences that can be modeled in the spatial normalization procedure and (ii) mesoscopic volume differences that cannot. The latter persist after normalization and are detected after spatial smoothing of grey matter maps (i.e., partitions or segments) – smoothing transforms these volume differences into image intensity differences through the partial volume effect.

More recently, this perspective has changed with the incorporation of an additional step, introduced to compensate for the effect of spatial normalization. When warping a series of images to match a template, it is inevitable that volumetric differences will be introduced into the warped images. For example, if one subject's temporal lobe has half the volume of that of the template, then its volume will be doubled during spatial normalization. This will also result in a doubling of the voxels labeled grey matter. To remove this confound, the spatially normalized grey matter (or other tissue class) is adjusted by multiplying by its relative volume before and after warping. If warping results in a region doubling its volume, then the correction will halve the intensity of the tissue label. This whole procedure has the effect of preserving the total amount of grey matter signal in the normalized partitions [19, 12] [6]. Classical VBM assumed that the warps were so smooth that these volume changes could be ignored. However, advances in normalization techniques now allow for high-resolution warps.

If brain images from different subjects can be matched together exactly [7], then a complete analysis of the volumetric differences could proceed using only the derived warps. Expansions and contractions occur when an image from one subject is warped to match that of another. These volume changes (hence also the relative volumes of structures) are encoded by the warps. One form of tensor-based morphometry (TBM) involves analyzing these relative volumes (more

[6] Note that a uniformly smaller brain will have uniformly lower grey matter intensities after the correction. Any detected differences are therefore less regionally specific, unless some kind of "global" measures are modeled as confounding effects during the statistical analyses. These could pertain to the total amount of grey matter in each brain, or, more usefully in many cases, the intra-cranial volume of each subject.

[7] Registered so that corresponding brain structures are matched, rather than solving the simpler problem of matching grey matter with grey matter, white with white and CSF with CSF.

formally, the Jacobian determinants of the deformation field) in order to identify regions of systematic volumetric difference. The adjustment step mentioned above can be considered from the perspective of TBM, in which volumetric changes derived from the warps are endowed with tissue specificity. By multiplying the relative volumes by the tissue class of interest, volumetric information about other tissue classes is discounted (i.e., there will be no changes attributed to grey matter in purely white matter regions). In other words, the product of grey matter and volume change has two equivalent interpretations: (i) in VBM it represents the proportion of the voxel that is grey matter, having adjusted for the confounding effects of warping the brains, and (ii) it represents the proportion of volume change attributable to grey matter.

By including the multiplication step, a continuum is introduced between the methods of VBM and TBM. At the extreme where little or no warping is used, most of the information pertaining to volumetric differences will be derived from systematic differences in the spatial distribution of the tissue under study. At the other extreme, where registration between images is exact, all the volumetric information will be encoded by the warps (because the normalized partitions will be identical). The product remains sensitive to differences at either extreme and can be regarded as an integration of classical VBM and TBM.

This dual perspective is illustrated in Figure 6.6, using brain images of a subject suffering from Alzheimers disease. This example illustrates progressive volumetric changes of CSF, particularly in the ventricles. The first column of the figure shows how the later of the two images was processed by classifying the CSF (see Chapter 5) and smoothing it using an isotropic 12mm FWHM Gaussian kernel. The column in the center shows the processing of the earlier image, which was first rigidly registered (see Chapter 2) with the late image and then the CSF was classified and smoothed. The difference between the smoothed CSF of the two images can be considered as analogous to VBM, where the spatial normalization is restricted to a rigid body registration.

The third column shows processing that is analogous to the augmented approach. After rigid registration, the early image is precisely warped to match the late one. The warping method (see Chapter 4) attempts to estimate exact displacements at every voxel, so is able to model the relative shapes of the pair of brains. The warped image is then classified, producing an image of CSF that is very similar to the CSF of the late image. A subtraction of these images would probably not show any meaningful differences. If the segmentation and warping were perfect, then the late CSF image would be identical to the warped early CSF image. To localize CSF volume differences, the volume changes resulting from the warping must also enter into the comparison. To do this, the warped CSF is simply multiplied by the relative volumes estimated from the warps. This means that the procedure preserves the amount of CSF from the original image, while also achieving a good registration. The figure shows an image of relative volumes, where lighter areas indicate a smaller volume in the early image. Following this multiplication, the data are smoothed. The difference between this image and the processed late image shows a picture of volumetric differences based on the warps (c.f., TBM). As can be seen from the bottom row of Figure 6.6, the pattern of CSF volume differences estimated using the two methods is very similar.

6.4.1 Method

The steps involved in VBM are:

- The procedure begins by spatially normalizing all subjects' data to the same stereotaxic space, which is achieved by registering each image to the same template. The methods described in Chapter 3 can be used to do this, although more precise methods can also be used. It is important that the quality of the registration is high, and that the choice of

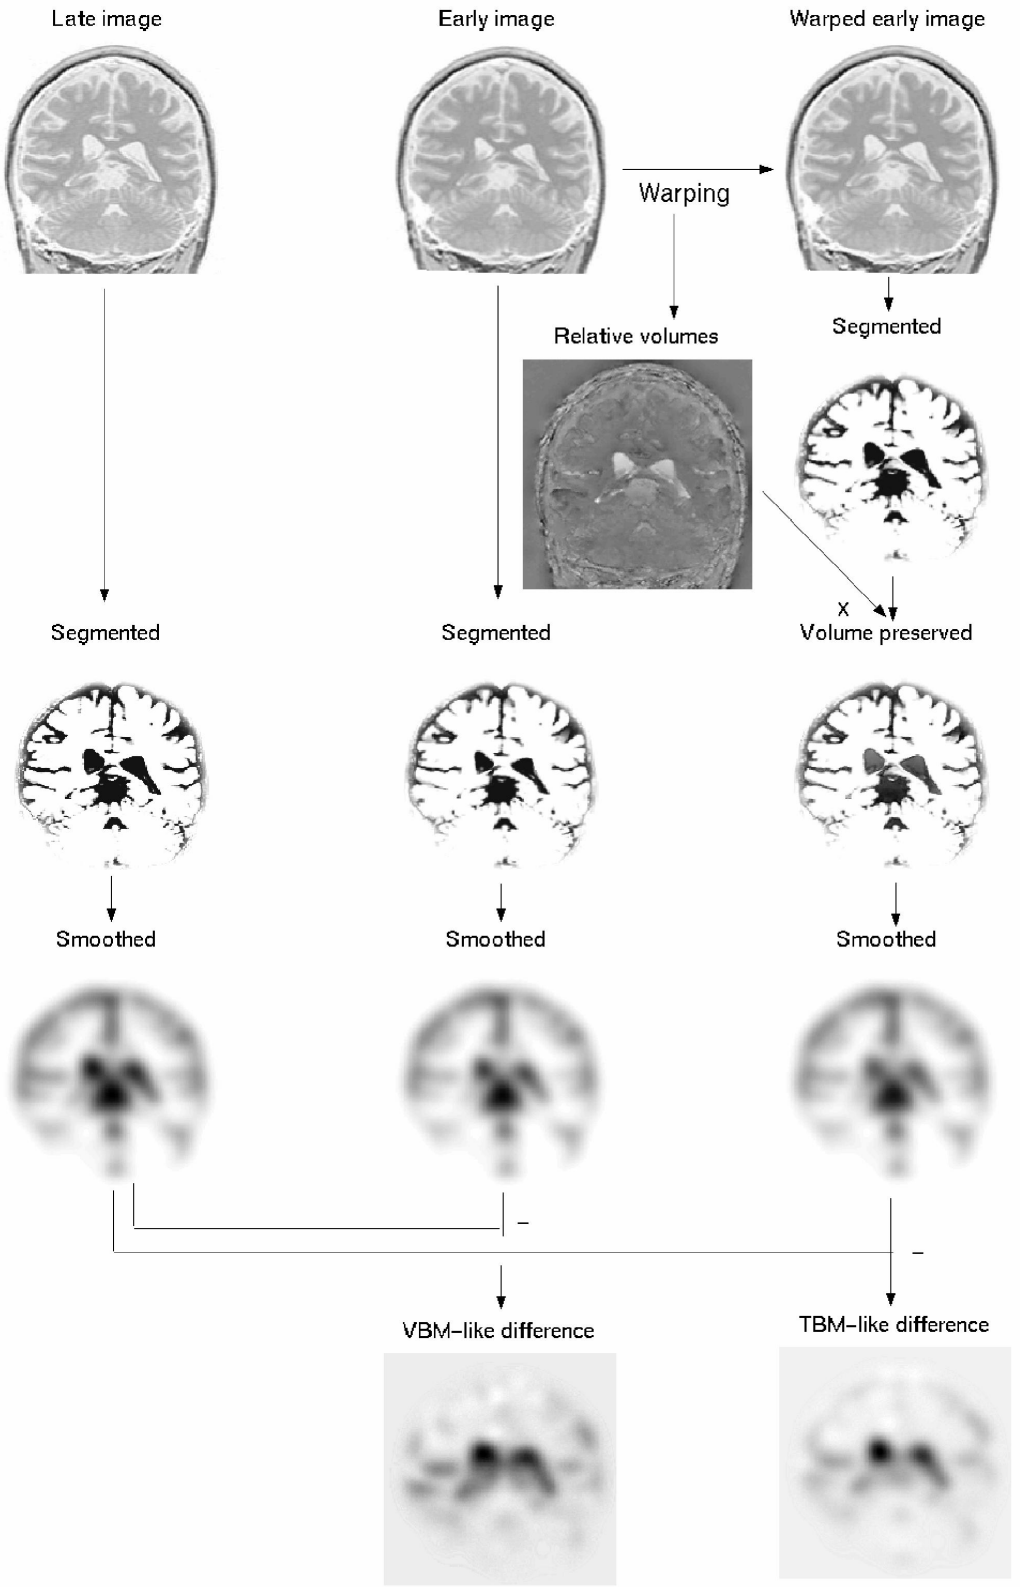

Figure 6.6: An illustration of the continuum between VBM and TBM. The center column illustrates a VBM processing stream, whereas the column on the right illustrates a TBM-like stream. The data were obtained from the Dementia Research Group, Queen Square, London.

template image does not bias the final solution. An ideal template consists of the average of a large number of MR images that have been registered to within the accuracy of the spatial normalization technique. The spatially normalized images should have a relatively high-resolution (1mm or 1.5mm isotropic voxels), so that the following grey matter extraction method is not excessively confounded by partial volume effects, where voxels contain a mixture of different tissue types.

The basis function approach of Chapter 3 is far from perfect [8], but, with the right data, it gives a good "global" match between brain images. Because of this, there are cases when structural differences, not directly related to grey matter volumes, can be identified as significant. One obvious example is when one population has larger ventricles. Because the spatial normalization method [1] cannot achieve an exact match, it has to change the volume of surrounding tissue when it attempts to make the ventricles of the individual subjects the same size. For example, if the ventricles are enlarged during spatial normalization, then grey matter near to them also also needs to be enlarged. It is then possible that structural differences pertaining to ventricular volume can show up in a VBM study of grey matter volume. A way of circumventing this would be to base the spatial normalization on only the segmented grey matter. If all the data entering into the statistical analysis were only derived from grey matter, then any significant differences must be due to grey matter.

- The spatially normalized images are then typically partitioned into different tissue classes using the segmentation technique described in Chapter 5. Note that if the segmentation step does not require prior probability maps of different tissue classes to be overlayed, then it should probably be performed first, in order to reduce the partial volume effects incurred by re-sampling the data. A further possible step after segmentation would be the binarization of the resulting tissue class images. Many tissue classification methods produces images where each voxel is the *a posteriori* probability that that voxel should be assigned to a particular tissue type according to the model. These probabilities are values between zero and one. Binarization would involve assigning each voxel to its most likely tissue class.

- Nonlinear spatial normalization results in the volumes of certain brain regions increasing, whereas others decrease [19, 12]. This has implications for the interpretation of what VBM actually tests. The objective of VBM is to identify regional differences in the amount of a particular tissue (grey or white matter). To preserve the actual amounts of grey matter within each structure, a further processing step can be incorporated that multiplies the partitioned images by the relative voxel volumes. These relative volumes are simply the Jacobian determinants of the deformation field (see above). Without this adjustment, VBM can be thought of as comparing the *relative concentration* of grey matter (i.e., the proportion of grey matter to other tissue types within a region). With the adjustment, VBM compares the *absolute amount* of grey matter in different regions.

- The grey matter images are now smoothed by convolving with an isotropic Gaussian kernel. This makes the subsequent voxel by voxel analysis comparable to a region of interest approach, because each voxel in the smoothed images contains the average amount of grey matter from around the voxel (where the region around the voxel is defined by the form of the smoothing kernel). This is often referred to as "grey matter density", but should not be confused with cell packing density measured cytoarchitectonically. By the central limit theorem, smoothing also has the effect of rendering the data more normally distributed, thus increasing the validity of parametric statistical tests. Whenever possible, the size of

[8] Brain warping methods are still in their infancy. Most warping methods involve making MAP-like estimates of deformations, making use of sub-optimal guesses for the likelihood and prior probability distributions. A single high-dimensional MAP estimate is not that useful when there are many other possible solutions with similar posterior probability, and estimating expectations over all MAP estimates is not feasible given the high-dimensionality of the parameter space.

the smoothing kernel should be comparable to the size of the expected regional differences between the groups of brains. The smoothing step also helps to compensate for the inexact nature of the spatial normalization.

- Following the pre-processing, which involves spatial normalization, tissue classification and spatial smoothing, the final step in a VBM analysis is to perform voxel-wise statistical tests. The results of these tests are a statistical parametric map (SPM) [16] showing significant regional differences among the populations included in the study. Corrections for multiple dependent comparisons are then made using the theory of Gaussian random fields [15, 25, 24].

  Statistical analysis using the general linear model (GLM) is used to identify regions of grey matter concentration that are significantly related to the particular effects under study [16]. The GLM is a flexible framework that allows many different tests to be applied, ranging from group comparisons and identification of regions of grey matter concentration that are related to specified covariates such as disease severity or age, to complex interactions between different effects of interest. Standard parametric statistical procedures (t-tests and F-tests) are used to test the hypotheses, so they are valid providing the residuals, after fitting the model, are independent and normally distributed. If the statistical model is appropriate there is no reason why the residuals should not be independent, but there are reasons why they may not be normally distributed. The original segmented images contain values between zero and one, where most of the values are very close to either of the extremes. Only by smoothing the segmented images does the behavior of the residuals become more normally distributed.

  Following application of the GLM, the significance of any differences is ascertained using GRF theory [15, 25]. A voxel-wise statistical parametric map (SPM) comprises the result of many statistical tests, so it is necessary to correct for multiple dependent comparisons.

  Classical statistical tests cannot be used to prove a hypothesis, only to reject a null hypothesis. Any significant differences that are detected could be explained by a number of different causes, which are not disambiguated by the statistical inference. In other words, these tests are not concerned with accepting a specific hypothesis about the cause of a difference, but involve rejecting a null hypothesis with a given certainty. When the null hypothesis has been rejected, it does not impute a particular explanation for the difference if there are many potential causes that could explain it. The pre-processing employed by VBM manipulates the data so that the ensuing tests are more sensitive to some causes than others. In particular, VBM has been devised to be sensitive to systematic differences in the volumes of grey matter structures, but significant results can also arise for other reasons. Attribution of what may be the cause of any detected difference generally requires a careful characterization of the parameter estimates after the inference has been made.

  An objection to VBM is that it is sensitive to systematic shape differences attributable to misregistration from the spatial normalization step [7]. This is one of a large number of potential systematic differences that can arise [3]. For example, a particular subject group may move more in the scanner, so the resulting images contain motion artifact. This motion may interact with the segmentation to produce systematic classification differences. Another group may have systematic differences in the relative intensity of grey matter voxels compared to white matter, or may have to be positioned differently in the scanner. All these reasons, plus others, may produce differences that are detectable by VBM. They are all real differences among the data, but may not necessarily be due to reductions in grey matter density.

6.5 Discussion

The morphometric methods developed here all attempt to automatically identify neuro-anatomical differences, either from features of spatially normalized images, or from deformation fields that encode information about image shapes. Rather than focusing on particular structures, one important aspect of these methods is that the entire brain can be examined in a balanced way.

The different forms of morphometry have their own advantages and disadvantages, and the optimum approach may depend on the types of structural difference expected among the data. Where there are global patterns of difference, then the global approaches (that do not produce a SPM) may be more powerful, as they can model covariances between shapes of different structures. In contrast, the SPM approaches may be preferable where discrete focal differences are expected. The main disadvantages of global testing methods (as opposed to SPM methods), is that there are generally far more variables for each brain than there are brains in a study. This implies that covariances between all variables cannot be computed, so some form of dimensionality reduction (or regularization) is required, leading to inevitable data loss.

A related issue is about how to normalize morphometric measures in order to produce the most useful results. This is particularly true for voxel-wise approaches, because relationships between other, more global, factors are not modeled. For example, consider a group of subjects with smaller hippocampi, but also smaller brains. Should the reduced hippocampi be considered significantly different if they correlate with a smaller brain? This size difference could also relate to lengths or volumes of temporal lobe, or to a whole number of other measurements. Consider another example. If the right temporal lobe is relatively larger than the left temporal lobe for a particular group of subjects, then should these differences be localized in the left or right lobes? These are arguments for using global multivariate approaches.

Another challenge concerns visualizing and communicating the results of morphometric tests. Three dimensional volumes are quite difficult to visualize, especially within the limited space of most journals. However, the results of morphometric tests are often vector or tensor fields. These are quite difficult to visualize in two dimensions, but in three dimensions the challenge becomes much worse. The more useful results of global morphometric methods are the canonical variates that characterize the differences, and these are often some form of three dimensional vector or tensor field. In comparison, differences localized by voxel-wise methods can be relatively easily presented as a SPM. Although the reasons may appear trivial, voxel-wise approaches will probably come to dominate because their results can be explained and presented much more easily.

When parametric statistical methods are used, statistical designs such as comparisons between single subjects and whole groups may not be as reliable as more balanced designs involving comparisons between groups of several subjects. It may be necessary to resort to non-parametric methods [20, 8] for such cases. The alternative may involve developing suitable transforms for the data that render them better behaved.

Other factors also influence the validity and sensitivity of the different morphometric approaches. In particular, the warping method, used to register the images to the same stereotaxic space, has a significant influence on the results. For DBM and TBM, the statistical tests are based entirely on the deformation fields produced by the warping methods. The warping methods described in this book are all nonlinear optimization procedures, and therefore can be susceptible to reaching local minima, and hence non-optimal solutions. These local minima have negative consequences for subsequent statistical tests, as the estimated shapes do not reflect the true shapes of anatomical structures. Other problems occur when warping brains containing severe pathologies. For example, if a brain contains features that are not present in the template image, then an accurate match can not be achieved. The effects of this mismatch may also propagate to other brain regions because of the inherent smoothness of the deformation fields. Much more

work is necessary in order to develop warping methods that can model the various forms of severe pathology that may be encountered.

Bibliography

[1] J. Ashburner and K. J. Friston. Nonlinear spatial normalization using basis functions. *Human Brain Mapping*, 7(4):254–266, 1999.

[2] J. Ashburner and K. J. Friston. Voxel-based morphometry – the methods. *NeuroImage*, 11:805–821, 2000.

[3] J. Ashburner and K. J. Friston. Why voxel-based morphometry should be used. *NeuroImage*, 14(6):1238–1243, 2001.

[4] F. L. Bookstein. Landmark methods for forms without landmarks: Morphometrics of group differences in outline shape. *Medical Image Analysis*, 1(3):225–243, 1997.

[5] F. L. Bookstein. Quadratic variation of deformations. In J. Duncan and G. Gindi, editors, *Proc. Information Processing in Medical Imaging*, pages 15–28, Berlin, Heidelberg, New York, 1997. Springer-Verlag.

[6] F. L. Bookstein. *Brain Warping*, chapter 10, pages 157–182. Academic Press, 1999.

[7] F. L. Bookstein. Voxel-based morphometry should not be used with imperfectly registered images. *NeuroImage*, 14(6):1454–1462, 2001.

[8] E. Bullmore, J. Suckling, S. Overmeyer, S. Rabe-Hesketh, E. Taylor, and M. Brammer. Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural MR images of the brain. *IEEE Transactions on Medical Imaging*, 18(1):32–42, 1999.

[9] J. Cao and K. J. Worsley. The geometry of the Hotelling's T^2 random field with applications to the detection of shape changes. *Annals of Statistics*, 27(3):925–942, 1999.

[10] M. K. Chung, K. J. Worsley, T. Paus, C. Cherif, D. L. Collins, J. N. Giedd, J. L. Rapoport, and A. C. Evans. A unified statistical approach to deformation-based morphometry. *NeuroImage*, 14(3):595–606, 2001.

[11] C. Davatzikos. *Handbook of Medical Imaging*, chapter 16, pages 249–260. Academic Press, 2000.

[12] C. Davatzikos, A. Genc, D. Xu, and S. M. Resnick. Voxel-based morphometry using the ravens maps: Methods and validation using simulated longitudinal atrophy. *NeuroImage*, 14(6):1361–1369, 2001.

[13] I. L. Dryden and K.V. Mardia. *Statistical Shape Analysis*. John Wiley and Sons, 1998.

[14] P. A. Freeborough and N. C. Fox. Modelling brain deformations in alzheimer disease by fluid registration of serial MR images. *Journal of Computer Assisted Tomography*, 22(5):838–843, 1998.

[15] K. J. Friston, A. P. Holmes, J.-B. Poline, C. J. Price, and C. D. Frith. Detecting activations in PET and fMRI: Levels of inference and power. *NeuroImage*, 4:223–235, 1996.

[16] K. J. Friston, A. P. Holmes, K. J. Worsley, J.-B. Poline, C. D. Frith, and R. S. J. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2:189–210, 1995.

[17] C. Gaser, H.-P. Volz, S. Kiebel, S. Riehemann, and H. Sauer. Detecting structural changes in whole brain based on nonlinear deformations - application to schizophrenia research. *NeuroImage*, 10:107–113, 1999.

[18] J. C. Gee and R. K. Bajcsy. *Brain Warping*, chapter 11, pages 183–198. Academic Press, 1999.

[19] A. F. Goldszal, C. Davatzikos, D. L. Pham, M. X. H. Yan, R. N. Bryan, and S. M. Resnick. An image-processing system for qualitative and quantitative volumetric analysis of brain images. *Journal of Computer Assisted Tomography*, 22(5):827–837, 1998.

[20] A. P. Holmes, R. C. Blair, D. G. Watson J, and I. Ford. Non-parametric analysis of statistic images from functional mapping experiments. *Journal of Cerebral Blood Flow and Metabolism*, 16:7–22, 1996.

[21] R. W. Ogden. *Non-linear Elastic Deformations*. Dover, 1984.

[22] P. M. Thompson, J. N. Giedd, R. P. Woods, D. MacDonald, A. C. Evans, and A. W. Toga. Growth patterns in the developing brain detected by using continuum mechanical tensor maps. *Nature*, 404:190–193, 2000.

[23] P. M. Thompson and A. W. Toga. *Brain Warping*, chapter 18, pages 311–336. Academic Press, 1999.

[24] K. J. Worsley, M. Andermann, T. Koulis, D. MacDonald, and A. C. Evans. Detecting changes in non-isotropic images. *Human Brain Mapping*, 8(2):98–101, 1999.

[25] K. J. Worsley, S. Marrett, P. Neelin, A. C. Vandal, K. J. Friston, and A. C. Evans. A unified statistical approach for determining significant voxels in images of cerebral activation. *Human Brain Mapping*, 4:58–73, 1996.

[26] I. C. Wright, Z. R. Ellison, T. Sharma, K. J. Friston, R. M. Murray, and P. K. Mcguire. Mapping of grey matter changes in schizophrenia. *Schizophrenia Research*, 35:1–14, 1999.

[27] I. C. Wright, P. K. McGuire, J.-B. Poline, J. M. Travere, R. M. Murray, C. D. Frith, R. S. J. Frackowiak, and K. J. Friston. A voxel-based method for the statistical analysis of gray and white matter density applied to schizophrenia. *NeuroImage*, 2:244–252, 1995.