Nonlinear Spatial Normalization Using Basis Functions

John Ashburner* and Karl J. Friston

*Functional Imaging Laboratory, Wellcome Department of Cognitive Neurology,
Institute of Neurology, London, United Kingdom*

Abstract: We describe a comprehensive framework for performing rapid and automatic nonlabel-based nonlinear spatial normalizations. The approach adopted minimizes the residual squared difference between an image and a template of the same modality. In order to reduce the number of parameters to be fitted, the nonlinear warps are described by a linear combination of low spatial frequency basis functions. The objective is to determine the optimum coefficients for each of the bases by minimizing the sum of squared differences between the image and template, while simultaneously maximizing the smoothness of the transformation using a *maximum a posteriori* (MAP) approach. Most MAP approaches assume that the variance associated with each voxel is already known and that there is no covariance between neighboring voxels. The approach described here attempts to estimate this variance from the data, and also corrects for the correlations between neighboring voxels. This makes the same approach suitable for the spatial normalization of both high-quality magnetic resonance images, and low-resolution noisy positron emission tomography images. A fast algorithm has been developed that utilizes 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. *Hum. Brain Mapping* 7:254–266, 1999.

© 1999 Wiley-Liss, Inc.

Key words: registration; anatomy; imaging; stereotaxy; basis functions; spatial normalization; PET; MRI; functional mapping

INTRODUCTION

This paper concerns the problem of nonlinear spatial normalization: namely, how to map a single subject's brain image into a standard space. The solution of this problem allows for a wide range of voxel-based analyses and facilitates the comparison of different subjects and data bases. The problem of spatial normalization

is not a trivial one; indeed, on some anatomical scales it is not clear that a solution even exists.

A fundamental advantage of using spatially normalized images is that activations can be reported according to a set of meaningful Euclidian coordinates within a standard space [Fox, 1995]. New results can be readily incorporated into ongoing brain atlas and data base projects such as that being developed by the International Consortium for Human Brain Mapping (ICBM) [Mazziotta et al., 1995]. The most commonly adopted coordinate system within the brain imaging community is that described by the atlas of Talairach and Tournoux [1988].

When whole-brain structural images (typically high-resolution magnetic resonance imaging; MRI) of the

Grant sponsor: Wellcome Trust.

*Correspondence to: John Ashburner, Functional Imaging Laboratory, Wellcome Department of Cognitive Neurology, Institute of Neurology, 12 Queen Square, London WC1N 3BG, UK. E-mail: j.ashburner@fil.ion.ucl.ac.uk

Received 21 October 1997; Accepted 25 January 1999

© 1999 Wiley-Liss, Inc.

subject are available in addition to the functional images, the images can be coregistered using any one of a number of methods for intermodality registration [Pelizzari et al., 1988; Woods et al., 1992; Studholme et al., 1995; Collignon et al., 1995; Ashburner and Friston, 1997]. This allows the spatial transformations that warp the images to the reference space to be determined from the structural images. These warps can then be applied to the functional images. Because there are only six rigid body parameters required to map between the structural and functional images, the coregistration parameters can be determined fairly accurately. The structural images should have higher spatial resolution, less noise, and more structural information than the functional images, allowing a more accurate nonlinear registration to be obtained.

However, not every functional imaging unit has ready access to a high-quality magnetic resonance (MR) scanner, so for many functional imaging studies there are no structural images of the subject available to the researcher. In this case, it is necessary to determine the required warps based solely on the functional images. These images may have a limited field of view, contain very little useful signal, or be particularly noisy. An ideal spatial normalization routine would need to be robust enough to cope with this type of data.

Nonlinear spatial transformations can be broadly divided into *label-based* and *nonlabel-based*. Label-based techniques identify homologous features (labels) in the image and template and find the transformations that best superpose them. The labels can be points, lines, or surfaces. Homologous features are often identified manually, but this process is time-consuming and subjective. Another disadvantage of using points as landmarks is that there are very few readily identifiable discrete points in the brain. A similar problem is faced during identification of homologous lines. However, surfaces are more readily identified, and in many instances they can be extracted automatically (or at least semiautomatically). Once they are identified, the spatial transformation is effected by bringing the homologies together. If the labels are points, then the required transformations at each of those points are known. Between the points, the deforming behavior is not known, so it is forced to be as “smooth” as possible. There are a number of methods for modeling this smoothness. The simplest models include fitting splines through the points in order to minimize *bending energy* [Bookstein, 1989]. More complex forms of interpolation are often used when the labels are sur-

faces. For example, Thompson and Toga [1996] mapped surfaces together using a fluid model.

Nonlabel-based approaches identify a spatial transformation that minimizes some index of the difference between an object and a template image, where both are treated as unlabeled continuous processes. The matching criterion is usually based upon minimizing the sum of squared differences or maximizing the correlation coefficient between the images. For this criterion to be successful, it requires the template to appear as a warped version of the image. In other words, there must be correspondence in the gray levels of the different tissue types between the image and template.

There are a number of approaches to nonlabel-based spatial normalization. A potentially enormous number of parameters are required to describe the nonlinear transformations that warp two images together (i.e., the problem is very high-dimensional). The forms of spatial normalization tend to differ in how they cope with the large number of parameters.

Some have abandoned conventional optimization approaches, and use viscous fluid models [Christensen et al., 1994, 1996] to describe the warps. In these models, finite-element methods are used to solve the partial differential equations that model one image as it “flows” to the same shape as the other. The major advantage of these methods is that they are able to account for large nonlinear displacements and also ensure that the topology of the warped image is preserved, but they do have the disadvantage that they are computationally expensive. Not every unit in the functional imaging field has the capacity to routinely perform spatial normalizations using these methods.

Others adopt a multiresolution approach whereby only a few of the parameters are determined at any one time [Collins et al., 1994b]. Usually, the entire volume is used to determine parameters that describe global low-frequency deformations. The volume is then subdivided, and slightly higher-frequency deformations are found for each subvolume. This continues until the desired deformation precision is achieved.

Another approach is to reduce the number of parameters that model the deformations. Some groups simply use only a 9- or 12-parameter affine transformation to spatially normalize their images, accounting for differences in position, orientation, and overall brain size. Low spatial frequency global variability in head shape can be accommodated by describing deformations by a linear combination of low-frequency basis functions [Amit et al., 1991]. The small number of parameters will not allow every feature to be matched

exactly, but it will permit the global head shape to be modeled. The method described in this paper is one such approach. The rationale for adopting a low-dimensional approach is that there is not necessarily a one-to-one mapping between any pair of brains. Different subjects have different patterns of gyral convolutions, and even if gyral anatomy can be matched exactly, this is no guarantee that areas of functional specialization will be matched in a homologous way. For the purpose of averaging signals from functional images of different subjects, very high-resolution spatial normalization may be unnecessary or unrealistic.

The deformations required to transform images to the same space are not clearly defined. Unlike rigid body transformations, where the constraints are explicit, those for nonlinear warping are more arbitrary. Without any constraints it is of course possible to transform any image such that it matches another exactly. The issue is therefore less about the nature of the transformation and more about defining constraints or priors under which a transformation is effected. The validity of a transformation can usually be reduced to the validity of these priors. Priors are normally incorporated using some form of Bayesian scheme, using estimators such as the *maximum a posteriori* (MAP) estimate or the *minimum variance estimate* (MVE). The MAP estimate is the single solution that has the highest a posteriori probability of being correct, and is the estimate that we attempt to obtain in this paper. The MVE was used by Miller et al. [1993], and is the solution that is the conditional mean of the posterior. The MVE is probably more appropriate than the MAP estimate for spatial normalization. However, if the errors associated with the parameter estimates and also the priors are normally distributed, then the MVE and the MAP estimates are identical.

The remainder of this paper is organized as follows. The theory section describes how the registrations are performed, beginning with a description of the Gauss-Newton optimization scheme employed. Following this, the paper describes specific implementational details for determining the optimal linear combination of spatial basis functions. This involves using properties of *Kronecker tensor products* for the rapid computation of the curvature matrix used by the optimization. A Bayesian framework using priors based upon membrane energy is then incorporated into the registration model. The following section provides an evaluation focusing on the utility of nonlinear deformations per se, and then the use of priors in a Bayesian framework.

There then follows a discussion of the issues raised, and of those for future consideration.

THEORY

There are two steps involved in registering any pair of images together. There is the *registration* itself, whereby the parameters describing a transformation are determined. Then there is the *transformation*, where one of the images is transformed according to the set of parameters. The registration step involves matching the object image to some form of standardized template image. Unlike in the work of Christensen et al. [1994; 1996] or the segmentation work of Collins et al. [1994a, 1995], spatial normalization requires that the images themselves are transformed to the space of the template, rather than a transformation being determined that transforms the template to the individual images.

The nonlinear spatial normalization approach described here assumes that the images have already been approximately registered with the template according to a 9 [Collins et al., 1994b] or 12-parameter [Ashburner et al., 1997] affine registration. This section will illustrate how the parameters describing global shape differences between the images and template are determined.

We begin by introducing a simple method of optimization based upon partial derivatives. Then the parameters describing the spatial transformations are introduced. In the current approach, the nonlinear warps are modeled by linear combinations of smooth basis functions, and a fast algorithm for determining the optimum combination of basis functions is described. For speed and simplicity, a relatively small number of parameters (approximately 1,000) are used to describe the nonlinear components of the registration.

The optimization method is extended to utilize Bayesian statistics in order to obtain a more robust fit. This requires knowledge of the errors associated with the parameter estimates, and also knowledge of the a priori distribution from which the parameters are drawn. This distribution is modeled in terms of a cost function based on the *membrane energy* of the deformations. We conclude with some comments on extending the schemes when matching images from different modalities.

Basic optimization algorithm

The objective of optimization is to determine a set of parameters for which some function is minimized (or

maximized). One of the simplest cases involves determining the optimum parameters for a model in order to minimize the sum of squared differences between the model and a set of real-world data (χ^2). Usually there are many parameters in the model, and it is not possible to exhaustively search through the whole parameter space. The usual approach is to make an initial estimate, and to iteratively search from there. At each iteration, the model is evaluated using the current estimates, and χ^2 is computed. A judgment is then made about how the parameters should be modified, before continuing on to the next iteration.

The image-registration approach described here is essentially an optimization. In the simplest case, one image (the object image) is spatially transformed so that it matches another (the template image), by minimizing χ^2 . In order for χ^2 minimization to work, it is important that a distorted version of the object image can be adequately modeled by the template (or templates). When using a single template image, it should have the same contrast distribution as the images that are to be matched to it. The parameters that are optimized are those that describe the spatial transformation (although there are often other nuisance parameters required by the model, such as intensity scaling parameters). The algorithm of choice [Friston et al., 1995] is one that is similar to *Gauss-Newton* optimization, and it is illustrated here:

Suppose that $e_i(p)$ is the function describing the difference between the object and template images at voxel i , when the vector of model parameters has values p . For each voxel (i), a first approximation of Taylor's theorem can be used to estimate the value that this difference will take if the parameters p are increased by t :

$$e_i(p + t) \approx e_i(p) + t_1 \frac{\partial e_i(p)}{\partial p_1} + t_2 \frac{\partial e_i(p)}{\partial p_2} \dots$$

This allows the construction of a set of simultaneous equations (of the form $Ax \approx e$) for estimating the values that t should assume in order to minimize $\sum_i e_i(p + t)^2$:

$$\begin{pmatrix} \frac{\partial e_1(p)}{\partial p_1} & \frac{\partial e_1(p)}{\partial p_2} & \dots \\ \frac{\partial e_2(p)}{\partial p_1} & \frac{\partial e_2(p)}{\partial p_2} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} t_1 \\ t_2 \\ \vdots \end{pmatrix} \approx \begin{pmatrix} e_1(p) \\ e_2(p) \\ \vdots \end{pmatrix}.$$

From this we can derive an iterative scheme for improving the parameter estimates. For iteration n , the parameters p are updated as:

$$p^{(n+1)} = p^{(n)} - (A^T A)^{-1} A^T e \quad (1)$$

where

$$A = \begin{pmatrix} \frac{\partial e_1(p)}{\partial p_1} & \frac{\partial e_1(p)}{\partial p_2} & \dots \\ \frac{\partial e_2(p)}{\partial p_1} & \frac{\partial e_2(p)}{\partial p_2} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} \text{ and } e = \begin{pmatrix} e_1(p) \\ e_2(p) \\ \vdots \end{pmatrix}.$$

This process is repeated until χ^2 can no longer be decreased, or for a fixed number of iterations. There is no guarantee that the best global solution will be reached, since the algorithm can get caught in a local minimum. The number of potential local minima is decreased by working with smooth images. This also has the effect of making the first-order Taylor approximation more accurate for larger displacements.

In practice, $A^T A$ and $A^T e$ from Eq. (1) are computed “on the fly” for each iteration. By computing these matrices using only a few rows of A and e at a time, much less computer memory is required than is necessary for storing the whole of matrix A . Also, the partial derivatives $\partial e_i(p)/\partial p_j$ can be rapidly computed from the gradients of the images using the chain rule.¹ These calculations will be illustrated more fully in the next section.

Parameterizing the spatial transformations

Spatial transformations are described by a linear combination of smooth basis functions. The choice of basis functions depends partly upon how translations at the boundaries should behave. If points at the boundary over which the transformation is computed are not required to move in any direction, then the basis functions should consist of the lowest frequencies of the three-dimensional discrete sine transform (DST). If there are to be no constraints at the boundaries, then a

¹Trilinear interpolation of the voxel lattice (rather than the sampling lattice) is used to resample the images at the desired coordinates. Gradients of the images are obtained at the same time, using a finite-difference method on the same lattice. No assumptions are made about voxel values that lie outside the field of view of image f . Points where y_i falls outside the domain of f are not included in the computations.

three-dimensional discrete cosine transform (DCT) is more appropriate. Both of these transforms use the same set of basis functions to represent warps in each of the directions. Alternatively, a mixture of DCT and DST basis functions can be used to constrain translations at the surfaces of the volume to be parallel to the surface only (*sliding* boundary conditions). By using a different combination of DCT and DST basis functions, the corners of the volume can be fixed and the remaining points on the surface can be free to move in all directions (*bending* boundary conditions) [Christensen, 1994]. The basis functions described here are the lowest-frequency components of the three (or two)-dimensional discrete cosine transform. In one dimension, the DCT of a function is generated by premultiplication with the matrix B^T , where the elements of the M by J matrix B are defined by:

$$b_{m,1} = \frac{1}{\sqrt{M}} \quad m = 1 \dots M$$

$$b_{m,j} = \sqrt{\frac{2}{M}} \cos\left(\frac{\pi(2m-1)(j-1)}{(2M)}\right)$$

$$m = 1 \dots M, j = 2 \dots J. \quad (2)$$

The two-dimensional DCT basis functions are shown in Figure 1, and a schematic example of a deformation based upon the DCT is shown in Figure 2.

The optimized parameters can be separated into a number of distinct groups. The most important are those for describing translations in the three orthogonal directions (t_1 , t_2 , and t_3). The model for defining the nonlinear warps uses deformations consisting of a linear combination of basis functions. In three dimensions, the transformation from coordinates x_i to coordinates y_i is:

$$y_{1,i} = x_{1,i} - u_{1,i} = x_{1,i} - \sum_{j=1}^J t_{j,1} b_{1,j}(x_i)$$

$$y_{2,i} = x_{2,i} - u_{2,i} = x_{2,i} - \sum_{j=1}^J t_{j,2} b_{2,j}(x_i)$$

$$y_{3,i} = x_{3,i} - u_{3,i} = x_{3,i} - \sum_{j=1}^J t_{j,3} b_{3,j}(x_i)$$

where t_{jd} is the j th coefficient for dimension d , and $b_{jd}(x)$ is the j th basis function at position x for dimension d .

Figure 1.

The lowest-frequency basis functions of a two-dimensional discrete cosine transform.

The optimization involves minimizing the sum of squared differences between the object image (f) and a template image (g). The images may be scaled differently, so an additional parameter (w) is needed to accommodate this difference. The minimized function is then:

$$\sum_i (f(y_i) - wg(x_i))^2.$$

Each element of vector e (from the previous section) contains $f(y_i) - wg(x_i)$. Derivatives of the function $f(y_i) - wg(x_i)$ with respect to each parameter are required in order to compute matrix A . These can be obtained using the chain rule:

$$\frac{\partial f(y_i)}{\partial t_{j,1}} = \frac{\partial f(y_i)}{\partial y_{1,i}} \frac{\partial y_{1,i}}{\partial t_{j,1}} = \frac{\partial f(y_i)}{\partial y_{1,i}} b_{j,1}(x_i)$$

$$\frac{\partial f(y_i)}{\partial t_{j,2}} = \frac{\partial f(y_i)}{\partial y_{2,i}} \frac{\partial y_{2,i}}{\partial t_{j,2}} = \frac{\partial f(y_i)}{\partial y_{2,i}} b_{j,2}(x_i)$$

$$\frac{\partial f(y_i)}{\partial t_{j,3}} = \frac{\partial f(y_i)}{\partial y_{3,i}} \frac{\partial y_{3,i}}{\partial t_{j,3}} = \frac{\partial f(y_i)}{\partial y_{3,i}} b_{j,3}(x_i).$$

In order to adopt the Gauss-Newton optimization strategy, $A^T A$ and $A^T e$ need to be computed on each iteration. Assuming that the lowest J frequencies of a three-dimensional DCT are used to define the warps, and there are I sampled positions in the image, then

Figure 2.

For the two-dimensional case, the deformation field consists of two scalar fields: one for horizontal deformations, and the other for vertical deformations. Images at left show the deformation fields as a linear combination of the basis images (see Fig. 1). Center column: Deformations in a more intuitive sense. The deformation field is applied by overlaying it on the object image, and resampling (right).

the theoretical size of the matrix A is $I \times (3J^3 + 1)$. The straightforward computation of $A^T A$ and $A^T e$ would be very time consuming. We now describe how this can be done in a much more expedient manner.

A fast algorithm

The fast algorithm for computing $A^T A$ and $A^T e$ is shown in Figure 3. The remainder of this section explains the matrix terminology used, and why it is so efficient.

For simplicity, the algorithm is only illustrated in two dimensions, although it has been implemented to estimate warps in three dimensions. Images f and g are considered as matrices F and G , respectively. Row m of F is denoted by $f_{m,:}$, and column n by $f_{:,n}$. The basis functions used by the algorithm are generated from a separable form from matrices B_1 and B_2 . By treating the transform coefficients as matrices T_1 and T_2 , the deformation

fields can be rapidly constructed by computing $B_1 T_1 B_2^T$ and $B_1 T_2 B_2^T$.

Between each iteration, image F is resampled according to the latest parameter estimates. The derivatives of F are also resampled to give $\nabla_1 F$ and $\nabla_2 F$. The i th element of each of these matrixes contain $f(y_i)$, $\partial f(y_i)/\partial y_{1i}$ and $\partial f(y_i)/\partial y_{2i}$, respectively.

The notation $\text{diag}(-\nabla_1 f_{:,m}) B_1$ simply means multiplying each element of row i of B_1 by $-\nabla_1 f_{i,m}$, and the symbol “ $\otimes$ ” refers to the *Kronecker tensor product*. If B_2 is a matrix of order $M \times J$, and B_1 is a second matrix, then:

$$B_2 \otimes B_1 = \begin{pmatrix} b_{211}B_1 & \cdots & b_{21J}B_1 \\ \vdots & \ddots & \vdots \\ b_{2M1}B_1 & \cdots & b_{2MJ}B_1 \end{pmatrix}.$$

The advantage of the algorithm shown in Figure 3 is that it utilizes some of the useful properties of Kro-

$$\alpha = \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \end{pmatrix}$$

for $m = 1 \dots M$

$$\begin{aligned} \mathbf{C} &= \mathbf{b}_{2m,:}^T \mathbf{b}_{2m,:} \\ \mathbf{E}_1 &= \text{diag}(-\nabla_1 \mathbf{f}_{:,m}) \mathbf{B}_1 \\ \mathbf{E}_2 &= \text{diag}(-\nabla_2 \mathbf{f}_{:,m}) \mathbf{B}_1 \end{aligned}$$

$$\alpha = \alpha + \begin{pmatrix} \mathbf{C} \otimes (\mathbf{E}_1^T \mathbf{E}_1) & \mathbf{C} \otimes (\mathbf{E}_1^T \mathbf{E}_2) & \mathbf{b}_{2m,:}^T \otimes (\mathbf{E}_1^T \mathbf{g}_{:,m}) \\ (\mathbf{C} \otimes (\mathbf{E}_1^T \mathbf{E}_2))^T & \mathbf{C} \otimes (\mathbf{E}_2^T \mathbf{E}_2) & \mathbf{b}_{2m,:}^T \otimes (\mathbf{E}_2^T \mathbf{g}_{:,m}) \\ (\mathbf{b}_{2m,:}^T \otimes (\mathbf{E}_1^T \mathbf{g}_{:,m}))^T & (\mathbf{b}_{2m,:}^T \otimes (\mathbf{E}_2^T \mathbf{g}_{:,m}))^T & \mathbf{g}_{:,m}^T \mathbf{g}_{:,m} \end{pmatrix}$$

$$\beta = \beta + \begin{pmatrix} \mathbf{b}_{2m,:}^T \otimes (\mathbf{E}_1^T (\mathbf{f}_{:,m} - w \mathbf{g}_{:,m})) \\ \mathbf{b}_{2m,:}^T \otimes (\mathbf{E}_2^T (\mathbf{f}_{:,m} - w \mathbf{g}_{:,m})) \\ \mathbf{g}_{:,m}^T (\mathbf{f}_{:,m} - w \mathbf{g}_{:,m}) \end{pmatrix}$$

end

Figure 3.

Two-dimensional illustration of the fast algorithm for computing $\mathbf{A}^T \mathbf{A}$ (α) and $\mathbf{A}^T \mathbf{e}$ (β).

necker tensor products. This is especially important when the algorithm is implemented in three dimensions. The performance enhancement results from a reordering of a set of operations such as $(\mathbf{B}_3 \otimes \mathbf{B}_2 \otimes \mathbf{B}_1)^T (\mathbf{B}_3 \otimes \mathbf{B}_2 \otimes \mathbf{B}_1)$, to the equivalent $(\mathbf{B}_3^T \mathbf{B}_3) \otimes (\mathbf{B}_2^T \mathbf{B}_2) \otimes (\mathbf{B}_1^T \mathbf{B}_1)$. Assuming that the matrices $\mathbf{B}_3$, $\mathbf{B}_2$, and $\mathbf{B}_1$ all have order $M \times J$, then the number of floating-point operations is reduced from $M^3 J^3 (J^3 + 2)$ to approximately $3M(J^2 + J) + J^6$. If M equals 32, and J equals 4, we expect a performance increase of about a factor of 20,000. The limiting factor to the algorithm is no longer the time taken to create the curvature matrix ($\mathbf{A}^T \mathbf{A}$), but is now the amount of memory required to store it, and the time taken to invert it.

A maximum a posteriori solution

Without regularization, it is possible to introduce unnecessary deformations that only reduce the residual sum of squares by a tiny amount. This could potentially make the algorithm very unstable. Regularization is achieved using Bayesian statistics.

Bayes' rule can be expressed in the continuous form as $p(\mathbf{a}_p | \mathbf{b}) \propto p(\mathbf{b} | \mathbf{a}_p) p(\mathbf{a}_p)$ where $p(\mathbf{a}_p)$ is the prior probability of $\mathbf{a}_p$ being true, $p(\mathbf{b} | \mathbf{a}_p)$ is the conditional probability that $\mathbf{b}$ is observed given that $\mathbf{a}_p$ is true, and

$p(\mathbf{a}_p | \mathbf{b})$ is the posterior probability of $\mathbf{a}_p$ being true, given that measurement $\mathbf{b}$ has been made. The maximum a posteriori (MAP) estimate for parameters $\mathbf{p}$ is the mode of $p(\mathbf{a}_p | \mathbf{b})$. For our purposes, $p(\mathbf{a}_p)$ represents a known prior probability distribution from which the parameters are drawn, $p(\mathbf{b} | \mathbf{a}_p)$ is the likelihood of obtaining the data $\mathbf{b}$ given the parameters (the maximum likelihood estimate), and $p(\mathbf{a}_p | \mathbf{b})$ is the function to be maximized.

A probability is related to its Gibb's form by $p(\mathbf{a}) \propto e^{-H(\mathbf{a})}$. Therefore, the posterior probability is maximized when its Gibb's form is minimized. This is equivalent to minimizing $H(\mathbf{b} | \mathbf{a}_p) + H(\mathbf{a}_p)$. In this expression, $H(\mathbf{b} | \mathbf{a}_p)$ is related to the residual sum of squares. If we assume that the parameters are drawn from a zero mean, multinormal distribution described by a covariance matrix $\mathbf{C}_o$, then $H(\mathbf{a}_p)$ is simply given by $\mathbf{p}^T \mathbf{C}_o \mathbf{p}$.

We previously illustrated a method of incorporating knowledge about a priori parameter distributions into the Gauss-Newton optimization scheme [Ashburner et al., 1997]. For a zero mean prior distribution, the iterative scheme is:

$$\mathbf{p}^{(n+1)} = (\mathbf{C}_o^{-1} \sigma^2 + \mathbf{A}^T \mathbf{A})^{-1} (\mathbf{A}^T \mathbf{A} \mathbf{p}^{(n)} - \mathbf{A}^T \mathbf{e}) \quad (3)$$

where:

$$\sigma^2 = \sum_{i=1}^I e_i(p)^2 / \nu. \quad (4)$$

ν refers to the number of degrees of freedom after correction for correlations between neighboring voxels.

The a priori distribution

The first requirement for a MAP approach is to define some form of prior distribution for the parameters. For a simple linear approach, the priors consist of an a priori estimate of the mean of the parameters (assumed to be zero), and also a covariance matrix describing the distribution of the parameters about this mean. There are many possible forms for modeling these priors, each of which refers to some type of “energy” term. The form of regularization described here is based upon the *membrane energy* or *laplacians* of the deformation field [Amit et al., 1991; Gee et al., 1997]. Two other types of linear regularization (*bending energy* and *linear-elastic energy*) are described in the Appendix. None of these schemes enforce a strict one-to-one mapping between the object and template images, but this makes little difference for the small deformations that we are interested in here.

In three dimensions, the membrane energy of the deformation field u is:

$$\sum_i \sum_{j=1}^3 \sum_{k=1}^3 \lambda \left(\frac{\partial u_{ji}}{\partial x_{ki}} \right)^2$$

where λ is simply a scaling constant. The membrane energy can be computed from the coefficients of the basis functions by $t_1^T H t_1 + t_2^T H t_2 + t_3^T H t_3$, where t_1 , t_2 , and t_3 refer to vectors containing the parameters describing translations in the three dimensions. The matrix H is defined by:

$$\begin{aligned} H = & \lambda (B_3^T B_3') \otimes (B_2^T B_2) \otimes (B_1^T B_1) \\ & + \lambda (B_3^T B_3) \otimes (B_2^T B_2') \otimes (B_1^T B_1) \\ & + \lambda (B_3^T B_3) \otimes (B_2^T B_2) \otimes (B_1'^T B_1') \end{aligned}$$

where the notation B_1' refers to the first derivatives of B_1 .

Assuming that the parameters consist of $(t_1^T t_2^T t_3^T w)^T$, matrix C_0^{-1} from Eq. (3) can be constructed from H by:

$$C_0^{-1} = \begin{pmatrix} H & 0 & 0 & 0 \\ 0 & H & 0 & 0 \\ 0 & 0 & H & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (5)$$

H is all zeros, except for the diagonal. Elements on the diagonal represent the reciprocal of the a priori variance of each parameter, and each is given by:

$$h_{j+J(k+J \times l)} = \lambda \pi^2 M^{-2} ((j-1)^2 + (k-1)^2 + (l-1)^2)$$

where M is the dimension of the DCT (see Eq. 2), and J is the number of low-frequency coefficients used in any dimension.

Values of λ that are too large will provide too much regularization and result in underestimated deformations. If the values are too small, there will not be enough regularization and the resulting deformations will include a large amount of noise. There is no simple way of estimating what the best values for these constants should be.

In the absence of known priors, the membrane energy provides a useful model in which we assume that the probability of a particular set of parameters arising is inversely related to the membrane energy associated with that set. Clearly this model is somewhat ad hoc, but it is a useful and sensible one. If the true prior distribution of the parameters is known (derived from a large number of subjects), then C_0 could be an empirically determined covariance matrix describing this distribution. This approach would have the advantage that the resulting deformations are more typically “brain-like,” and so increase the face validity of the approach.

Templates and intensity transformations

So far, only a single intensity scaling parameter (w) has been considered. This is most effective when there is a linear relationship between the images. However, for spatially normalizing some images, it is necessary to include other parameters describing intensity transformations.

The optimization can be assumed to minimize two sets of parameters: those that describe spatial transformations (p_i), and those for describing intensity transformations (p_w). This means that the difference function

can be expressed in the generic form:

$$e_i(p) = f(t(x_i, p_t)) - w(x_i, p_w)$$

where f is the object image, $t()$ is a vector function describing the spatial transformations based upon parameters p_t , and $w()$ is a scalar function describing intensity transformations based on parameters p_w . x_i represents the coordinates of the i th sampled point.

The intensities could vary spatially (e.g., due to inhomogeneities in the MRI scanner). Linear variations in intensity over the field of view can be accounted for by optimizing a function of the form:

$$\sum_i (f(x_i, p_t) - (p_{w1}g(x_i) + p_{w2}x_{1i}g(x_i) + p_{w3}x_{2i}g(x_i) + p_{w4}x_{3i}g(x_i)))^2.$$

More complex variations could be included by modulating with other basis functions (such as the DCT basis function set described earlier) [Friston et al., 1995]. Information on the smoothness of the inhomogeneities could be incorporated by appropriate modifications to the matrix C_0^{-1} .

Another important idea is that a given image can be matched not to one reference image, but to a series of images that all conform to the same space. The idea here is that (ignoring the spatial differences) any given image can be expressed as a linear combination of a set of reference images. For example, these reference images might include different modalities (e.g., positron emission tomography (PET), single photon emission computed tomography (SPECT), ^{18}F -DOPA, ^{18}F -deoxy-glucose, T_1 -weighted MRI, and T_2^* -weighted MRI) or different anatomical tissues (e.g., gray matter, white matter, and cerebro-spinal fluid (CSF) segmented from the same T_1 -weighted MRI) or different anatomical regions (e.g., cortical gray matter, subcortical gray matter, and cerebellum) or finally any combination of the above. Any given image, irrespective of its modality, could be approximated with a function of these images. A simple example using two images would be:

$$\sum_i (f(Mx_i) - (p_{w1}g_1(x_i) + p_{w2}g_2(x_i)))^2.$$

Again, some form of model describing the likely a priori distributions of the parameters could be included.

EVALUATION

This section provides an anecdotal evaluation of the techniques presented in the previous section. We com-

pare spatial normalization both with and without nonlinear deformations, and compare nonlinear deformations with and without the use of Bayesian priors.

T_1 -weighted MR images of 12 subjects were spatially normalized to the same anatomical space. The normalizations were performed twice, first using only 12-parameter affine transformations and then using affine transformations followed by nonlinear transformations. The nonlinear transformation used 392 ($7 \times 8 \times 7$) parameters to describe deformations in each of the directions, and four parameters to model a linear scaling and simple linear image intensity inhomogeneities (making a total of 1,180 parameters in all). The basis functions were those of a three-dimensional DCT, and the regularization minimized the membrane energy of the deformation fields (using a value of 0.01 for λ). Twelve iterations of the nonlinear registration algorithm were performed.

Figure 4 shows pixel-by-pixel means and standard deviations of the normalized images. The mean image from the nonlinear normalization shows more contrast and has edges that are slightly sharper. The standard deviation image for the nonlinear normalization shows decreased intensities, demonstrating that the intensity differences between the images have been reduced. However, the differences tend to reflect changes in the global shape of the heads, rather than differences between cortical anatomy.

This evaluation should illustrate the fact that the nonlinear normalization clearly reduces the sum of squared intensity differences between the images. The amount of residual variance could have been reduced further by decreasing the amount of regularization. However, this may lead to some very unnatural-looking distortions being introduced, due to an overestimation of the a priori variability. Evaluations like this tend to show more favorable results for less heavily regularized algorithms. With less regularization, the optimum solution is based more upon minimizing the difference between the images, and less upon knowledge of the a priori distribution of the parameters. This is illustrated for a single subject in Figure 5, where the distortions of gyral anatomy clearly have a very low face validity (Fig. 5, lower right).

DISCUSSION

The criteria for “good” spatial transformations can be framed in terms of validity, reliability, and computational efficiency. The validity of a particular transformation device is not easy to define or measure and indeed varies with the application. For example, a rigid body

Mean Images

Standard Deviation Images

Figure 4.

Means and standard deviations of spatially normalized T_1 -weighted images from 12 subjects. Images at left were derived using only affine registration. Those at right used nonlinear registration in addition to affine registration.

transformation may be perfectly valid for realignment but not for spatial normalization of an arbitrary brain into a standard stereotactic space. Generally the sorts of validity that are important in spatial transformations can be divided into (1) *face validity*, established by demonstrating that the transformation does what it is supposed to do, and (2) *construct validity*, assessed by comparison with other techniques or constructs. Face validity is a complex issue in functional mapping. At first glance, face validity might be equated with the coregistration of anatomical homologues in two images. This would be complete and appropriate if the biological question referred to structural differences or modes of variation. In other circumstances, however,

Figure 5.

Top left: Object or template image. Top right: An image that has been registered with it using a 12-parameter affine registration. Bottom left: Same image registered using the 12-parameter affine registration, followed by a regularized global nonlinear registration (using 1,180 parameters, 12 iterations, and a λ of 0.01). It should be clear that the shape of the image approaches that of the template much better after nonlinear registration. Bottom right: Image after the same affine transformation and nonlinear registration, but this time without using any regularization. The mean squared difference between the image and template after the affine registration was 472.1. After regularized nonlinear registration, this was reduced to 302.7. Without regularization, a mean squared difference of 287.3 was achieved, but this was at the expense of introducing a lot of unnecessary warping.

this definition of face validity is not appropriate. For example, the purpose of spatial normalization (either within or between subjects) in functional mapping studies is to maximize the sensitivity to neurophysiological change elicited by experimental manipulation of a sensorimotor or cognitive state. In this case, a better definition of a valid normalization is that which maximizes condition-dependent effects with respect to error (and if relevant, intersubject) effects. This will probably be effected when functional anatomy is congruent. This may or may not be the same as registering structural anatomy.

Limitations of nonlinear registration

Because deformations are only defined by a few hundred parameters, the nonlinear registration method described here does not have the potential precision of some other methods. High-frequency deformations cannot be modeled, since the deformations are restricted to the lowest spatial frequencies of the basis functions. This means that the current approach is unsuitable for attempting exact matches between fine cortical structures.

The method is relatively fast, taking on the order of 30 sec per iteration, depending upon the number of basis functions used. The speed is partly a result of the small number of parameters involved, and the simple optimization algorithm that assumes an almost quadratic error surface. Because the images are first matched using a simple affine transformation, there is less “work” for the algorithm to do, and a good registration can be achieved with only a few iterations (about 20). The method does not rigorously enforce a one-to-one match between the brains being registered. However, by estimating only the lowest-frequency deformations and by using appropriate regularization, this constraint is rarely broken.

When higher spatial frequency warps are to be fitted, more DCT coefficients are required to describe the deformations. There are practical problems that occur when more than about the $8 \times 8 \times 8$ lowest-frequency DCT components are used. One problem is that of storing and inverting the curvature matrix ($A^T A$). Even with deformations limited to $8 \times 8 \times 8$ coefficients, there are at least 1,537 unknown parameters, requiring a curvature matrix of about 18 M bytes (using double-precision floating-point arithmetic). Other methods which search for more parameters should be used when more precision is required. These include the method of Collins et al. [1994a], the high-dimensional linear-elasticity model [Miller et al., 1993], and the viscous fluid models [Christensen et al., 1996; Thompson and Toga, 1996].

In practice, however, it may be meaningless to even attempt an exact match between brains beyond a certain resolution. There is not a one-to-one relationship between the cortical structures of one brain and those of another, so any method that attempts to match brains exactly must be folding the brain to create sulci and gyri that do not really exist. Even if an exact match is possible, because the registration problem is not convex, the solutions obtained by high-dimensional warping techniques may not be truly optimum. These methods are very good at registering gray matter with

gray matter (for example), but there is no guarantee that the registered gray matter arises from homologous cortical features.

Also, structure and function are not always tightly linked. Even if structurally equivalent regions can be brought into exact register, it does not mean that the same is true for regions that perform the same or similar functions. For intersubject averaging, an assumption is made that functionally equivalent regions lie in approximately the same parts of the brain. This led to the current rationale for smoothing images from multisubject studies prior to performing analyses. Constructive interference of the smeared activation signals then has the effect of producing a signal that is roughly in an average location. In order to account for substantial fine-scale warps in a spatial normalization, it is necessary for some voxels to increase their volumes considerably, and for others to shrink to an almost negligible size. The contribution of the shrunken regions to the smoothed images is tiny, and the sensitivity of the tests for detecting activations in these regions is reduced. This is another argument in favor of registering only on a global scale.

The constrained normalization described here assumes that the template resembles a warped version of the image. Modifications are required in order to apply the method to diseased or lesioned brains. One possible approach is to assume different weights for different brain regions. Lesioned areas could be assigned lower weights, so that they would have much less influence on the final solution.

CONCLUSIONS

Consider the deformation fields required to map brain images to a common stereotactic space. Fourier-transforming the fields reveals that most of the variance is low-frequency, even when the deformations have been determined using good high-dimensional nonlinear registration methods. Therefore, an efficient representation of the fields can be obtained from the low-frequency coefficients of the transform. The current approach to spatial normalization utilizes this compression. Rather than estimating warps based upon literally millions of parameters, only a few hundred parameters are used to represent the deformations as a linear combination of a few low-frequency basis functions.

The method we have developed is automatic and nonlabel-based. A maximum a posteriori (MAP) approach is used to regularize the optimization. However, the main difference between this and other MAP

approaches is that an estimate of the errors is derived from the data per se. This estimate also includes a correction for local correlations between voxels. An implication of this is that the approach is suitable for spatially normalizing a wide range of different image modalities. High-quality MR images, and also low-resolution noisy PET images, can be treated the same way.

The spatial normalization converges rapidly, because it uses an optimization strategy with fast local convergence properties. Each iteration of the algorithm requires the computation of a Hessian matrix ($A^T A$). The straightforward computation of this matrix would be prohibitively time-consuming. However, this problem has been solved by developing an extremely fast method of computing this matrix that relies on the separable properties of the basis functions. A performance increase of several orders of magnitude is achieved in this way.

REFERENCES

- Amit Y, Grenander U, Piccioni M. 1991. Structural image restoration through deformable templates. *J Am Stat Assoc* 86:376–387.
- Ashburner J, Friston KJ. 1997. Multimodal image coregistration and partitioning—a unified framework. *Neuroimage* 6:209–217.
- Ashburner J, Neelin P, Collins DL, Evans AC, Friston KJ. 1997. Incorporating prior knowledge into image registration. *Neuroimage* 6:344–352.
- Bookstein FL. 1989. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Trans Pattern Anal Machine Intelligence* 11:567–585.
- Bookstein FL. 1997a. Landmark methods for forms without landmarks: morphometrics of group differences in outline shape. *Med Image Anal* 1:225–243.
- Bookstein FL. 1997b. Quadratic variation of deformations. In: Duncan J, Gindi G, editors. *Information processing in medical imaging*. New York: Springer-Verlag. p 15–28.
- Christensen GE. 1994. Deformable shape models for anatomy. Doctoral thesis. Sener Institute, Washington University.
- Christensen GE, Rabbitt RD, Miller MI. 1994. 3D brain mapping using a deformable neuroanatomy. *Phys Med Biol* 39:609–618.
- Christensen GE, Rabbitt RD, Miller MI. 1996. Deformable templates using large deformation kinematics. *IEEE Trans Image Process* 5:1435–1447.
- Collignon A, Maes F, Delaere D, Vandermeulen D, Suetens P, Marchal G. 1995. Automated multi-modality image registration based on information theory. In: Bizais Y, Barillot C, Di Paola, editors. *Information processing in medical imaging*. Dordrecht: Kluwer Academic Publishers. p 263–274.
- Collins DL, Peters TM, Evans AC. 1994a. An automated 3D nonlinear image deformation procedure for determination of gross morphometric variability in human brain. In: *Proc. Conference on Visualisation in Biomedical Computing*. p 180–190.
- Collins DL, Neelin P, Peters TM, Evans AC. 1994b. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *J Comput Assist Tomogr* 18:192–205.
- Collins DL, Evans AC, Holmes C, Peters TM. 1995. Automatic 3D segmentation of neuro-anatomical structures from MRI. In: Bizais Y, Barillot C, Di Paola R, editors. *Information processing in medical imaging*. Dordrecht: Kluwer Academic Publishers. p 139–152.
- Fox PT. 1995. Spatial normalization origins: objectives, applications, and alternatives. *Hum Brain Mapp* 3:161–164.
- Friston KJ, Ashburner J, Frith CD, Poline J-B, Heather JD, Frackowiak RSJ. 1995. Spatial registration and normalization of images. *Hum Brain Mapp* 2:165–189.
- Gee JC, Haynor DR, Le Briquer L, Bajcsy RK. 1997. Advances in elastic matching theory and its implementation. In: Goos G, Hartmanis J, Van Leeuwen J. *CVRMed-MRCAS'97*. New York: Springer-Verlag.
- Mazziotta JC, Toga AW, Evans A, Fox P, Lancaster J. 1995. A probabilistic atlas of the human brain: theory and rationale for its development. *Neuroimage* 2:89–101.
- Miller MI, Christensen GE, Amit Y, Grenander U. 1993. Mathematical textbook of deformable neuroanatomies. *Proc Natl Acad Sci USA* 90:11944–11948.
- Pelizzari CA, Chen GTY, Spelbring DR, Weichselbaum RR, Chen CT. 1988. Accurate three-dimensional registration of CT, PET and MR images of the brain. *J Comput Assist Tomogr* 13:20–26.
- Studholme C, Hill DLG, Hawkes DJ. 1995. Multiresolution voxel similarity measures for MR-PET coregistration. In: Bizais Y, Barillot C, Di Paola R. *Information Processing in Medical Imaging*. Dordrecht: Kluwer Academic Publishers. p 287–298.
- Talairach J, Tournoux. 1988. *Coplanar stereotaxic atlas of the human brain*. New York: Thieme Medical.
- Thompson PM, Toga AW. 1996. Visualization and mapping of anatomic abnormalities using a probabilistic brain atlas based on random fluid transformations. In: Hohne KH, Kikinis R, editors. *Proceedings of the International Conference on Visualization in Biomedical Computing*. New York: Springer-Verlag. pp 383–392.
- Woods RP, Cherry SR, Mazziotta JC. 1992. Rapid automated algorithm for aligning and reslicing PET images. *J Comput Assist Tomogr* 16:620–633.

OTHER LINEAR PRIORS

As an alternative to the *membrane energy* prior already discussed, we now show how two other priors can easily be incorporated into the model. For simplicity, they will only be shown for the two-dimensional case.

Bending energy

Bookstein's thin-plate splines [1997a,b] minimize the *bending energy* of the deformations. For the two-dimensional case, the bending energy of the deformation field is defined by:

$$\lambda \sum_i \left(\left(\frac{\partial^2 u_{1i}}{\partial x_{1i}^2} \right)^2 + \left(\frac{\partial^2 u_{1i}}{\partial x_{2i}^2} \right)^2 + 2 \left(\frac{\partial^2 u_{1i}}{\partial x_{1i} \partial x_{2i}} \right)^2 \right) + \lambda \sum_i \left(\left(\frac{\partial^2 u_{2i}}{\partial x_{1i}^2} \right)^2 + \left(\frac{\partial^2 u_{2i}}{\partial x_{2i}^2} \right)^2 + 2 \left(\frac{\partial^2 u_{2i}}{\partial x_{1i} \partial x_{2i}} \right)^2 \right).$$

This can be computed by:

$$\begin{aligned} & \lambda t_1^T (B_2'' \otimes B_1)^T (B_2'' \otimes B_1) t_1 + \lambda t_1^T (B_2 \otimes B_1'')^T (B_2 \otimes B_1') t_1 + \\ & 2\lambda t_1^T (B_2' \otimes B_1')^T (B_2' \otimes B_1') t_1 + \lambda t_2^T (B_2'' \otimes B_1)^T (B_2'' \otimes B_1) t_2 + \\ & \lambda t_2^T (B_2 \otimes B_1'')^T (B_2 \otimes B_1') t_2 + 2\lambda t_2^T (B_2' \otimes B_1')^T (B_2' \otimes B_1') t_2 \end{aligned}$$

where the notation B_1' and B_1'' refer to the first and second derivatives of B_1 . This is simplified to $t_1^T H t_1 + t_2^T H t_2$ where:

$$\begin{aligned} H = \lambda & ((B_2''^T B_2'') \otimes (B_1^T B_1) + (B_2^T B_2) \otimes (B_1'^T B_1')) \\ & + 2 (B_2'^T B_2') \otimes (B_1'^T B_1')). \end{aligned}$$

Matrix C_o^{-1} from Eq. (3) can be constructed from H as described by Eq. (5), but this time values on the diagonals are given by:

$$\begin{aligned} h_{j+k \times j} = \lambda & \left(\left(\frac{\pi(j-1)}{M} \right)^4 + \left(\frac{\pi(k-1)}{M} \right)^4 \right. \\ & \left. + 2 \left(\frac{\pi(j-1)}{M} \right)^2 \left(\frac{\pi(k-1)}{M} \right)^2 \right). \end{aligned}$$

Linear elasticity

The *linear-elastic* energy [Miller et al., 1993] of a two-dimensional deformation field is:

$$\sum_{j=1}^2 \sum_{k=1}^2 \sum_i \frac{\lambda}{2} \left(\frac{\partial u_{j,i}}{\partial x_{j,i}} \right) \left(\frac{\partial u_{k,i}}{\partial x_{k,i}} \right) + \frac{\mu}{4} \left(\frac{\partial u_{j,i}}{\partial x_{k,i}} + \frac{\partial u_{k,i}}{\partial x_{j,i}} \right)^2$$

where λ and μ are the *Lamé* elasticity constants. The elastic energy of the deformations can be computed by:

$$\begin{aligned} & (\mu + \lambda/2) t_1^T (B_2 \otimes B_1')^T (B_2 \otimes B_1') t_1 \\ & + (\mu + \lambda/2) t_2^T (B_2' \otimes B_1)^T (B_2' \otimes B_1) t_2 \\ & + \mu/2 t_1^T (B_2' \otimes B_1)^T (B_2' \otimes B_1) t_1 \\ & + \mu/2 t_2^T (B_2 \otimes B_1')^T (B_2 \otimes B_1') t_2 \\ & + \mu/2 t_1^T (B_2' \otimes B_1)^T (B_2 \otimes B_1') t_2 \\ & + \mu/2 t_2^T (B_2 \otimes B_1')^T (B_2' \otimes B_1) t_1 \\ & + \lambda/2 t_1^T (B_2 \otimes B_1')^T (B_2' \otimes B_1) t_2 \\ & + \lambda/2 t_2^T (B_2' \otimes B_1)^T (B_2 \otimes B_1') t_1. \end{aligned}$$

A regularization based upon this model requires an inverse covariance matrix (C_o^{-1}) that is not a simple diagonal matrix. This matrix is constructed as follows:

$$C_o^{-1} = \begin{pmatrix} H_1 & H_3 & 0 \\ H_3^T & H_2 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

where

$$H_1 = (\mu + \lambda/2) (B_2^T B_2) \otimes (B_1'^T B_1') + \mu/2 (B_2'^T B_2') \otimes (B_1^T B_1)$$

$$H_2 = (\mu + \lambda/2) (B_2'^T B_2') \otimes (B_1^T B_1) + \mu/2 (B_2^T B_2) \otimes (B_1'^T B_1')$$

$$H_3 = \lambda/2 (B_2^T B_2') \otimes (B_1'^T B_1) + \mu/2 (B_2'^T B_2) \otimes (B_1^T B_1').$$