On matching brain volumes*

J. C. Gee

Department of Neurology, University of Pennsylvania, Philadelphia 19104, USA

Abstract

To characterize the complex morphological variations that occur naturally in human neuroanatomy so that their confounding effect can be minimized in the identification of brain structures in medical images, a computational framework has evolved in which individual anatomies are modeled as warped versions of a canonical representation of the anatomy, known as an *atlas*. To realize this framework, the method of *elastic matching* was invented for determining the spatial mapping between a three-dimensional image pair in which one image volume is modeled as an elastic continuum that is deformed to match the appearance of the second volume. In this paper, we review the seminal ideas underlying the elastic matching technique, consider the practical implications of an *integral formulation* of the approach, and explore a more general *Bayesian* interpretation of the method in order to address issues that are less naturally resolved within a continuum mechanical setting, such as the examination of a solution's reliability or the incorporation of empirical information that may be available about the spatial mappings into the analysis.

1 Introduction

Relying on the knowledge that different instances of the same anatomy share a common topological scheme, Bajcsy and Broit proposed that the various problems of localizing, measuring, or visualizing anatomic structures in a computed tomographic study could be reduced to the single problem of determining the spatial mapping between the anatomy depicted in the study and a reference, or *atlas*, image of the same anatomy [1,2]. Using this mapping, information encoded in the atlas—such as anatomic labels and structural boundaries—is transferred to the image study. The transformed atlas labels automatically localize their corresponding structures in the tomographic images, as demonstrated in figure 1. Visualization is achieved by deforming, according to the

* This work was supported by the U.S.P.H.S. under grant RO1-NS-33662.

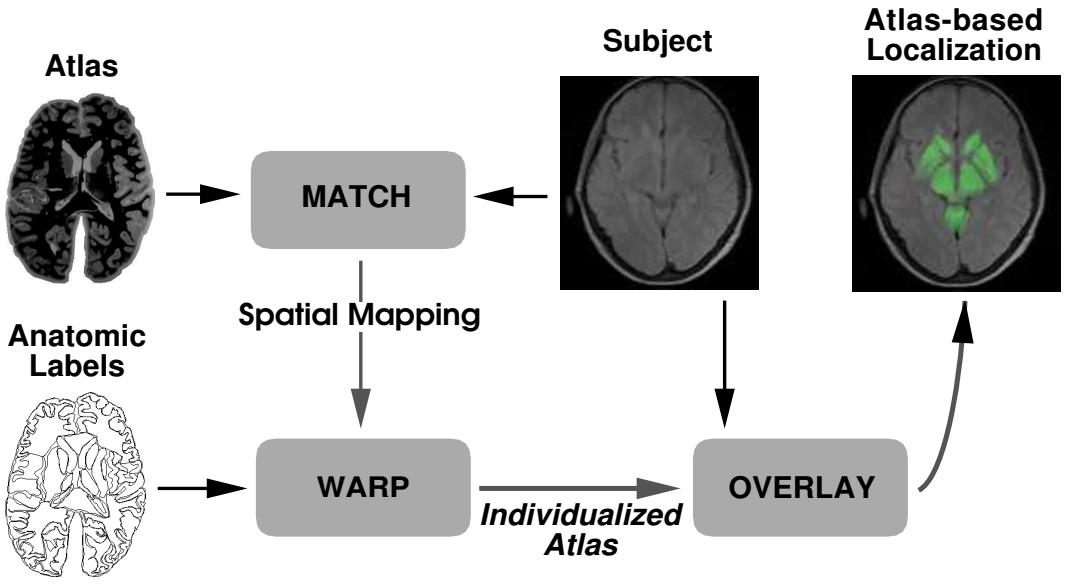

Fig. 1. Atlas-based approach to anatomic localization.

mapping, the structural boundaries contained in the atlas and then displaying the result. Quantitative measurements of any structure in the tomographic study can be inferred from the morphometric information associated with its warped atlas version.

To realize this method of analyzing tomographic images, Broit invented a general registration technique in which one image volume is modeled as an elastic continuum that is deformed to match the appearance of the second volume. In this paper, we review the seminal ideas underlying Broit's elastic matching method and then discuss their generalization within a probabilistic framework. We introduce the probabilistic approach by first studying the integral formulation of elastic matching. The integral form reveals important possibilities, including the use of powerful numerical solution methods and the existence of a Bayesian decision-theoretic interpretation of the registration problem.

The Bayesian paradigm associates with each model component of the problem a probability distribution for the model parameters, which quantifies the reliability of their values, and propagates this information to bear on subsequent analyses. As a result, the approach can potentially resolve a number of issues that are less naturally addressed within a continuum mechanical setting. These issues include, for example, the exploration of a solution's reliability or the incorporation of empirical information that may be available about the spatial transformations into an analysis. For computational anatomy [3–9], these aspects of the probabilistic approach will figure importantly in the development of a comprehensive methodology.

2 Elastic Matching

In conceiving the elastic matching, Broit sought a method that would optimally balance the similarity induced between the registered image pair with the amount of deformation, to the transformed image, required to achieve the result. Specifically, the solutions were made to optimize the objective function

$$\text{cost} = \text{deformation} - \text{similarity}. \tag{1}$$

Broit noted that erroneous correspondences could be established between two images when there was noise in the images or when certain portions of an image lacked the necessary features to distinguish them from neighboring regions. To help compensate for these errors, the mappings were presumed to be smooth, which explains the role of the deformation penalty in the problem formulation.

Broit drew an analogy between the elastic matching of two images and the physical process of applying forces to an elastic version of the object depicted in one of the images so that its deformed configuration resembles the target object in the second image. The appropriate forces are derived from a potential function—the *similarity* term in (1)—that provides at each point of the object under deformation a measure of the similarity between that point and any of the points in the target object. The first object is deformed in this way until an equilibrium configuration is reached or, equivalently, until the total potential energy of the system is at a local minimum [10].

Note that if the elastic strain energy is used as the measure of deformation in (1), the cost would be identical to the forementioned total potential energy. Therefore, the elastostatic configurations, which locally minimize this cost, represent solutions to the corresponding elastic matching problem.

2.1 Integral Formulation

Let σ and ε denote the stress and strain tensors, respectively [10]. Their elastic constitutive relation is given by the generalized Hooke's law: in indicial notation,

$$\sigma_{ij} = C_{ijkl}\,\varepsilon_{kl}, \tag{2}$$

where the elastic coefficients $C_{ijkl}(x)$ reduce to the two constants known as the Lamé parameters for homogeneous isotropic materials. Suppose further that u describes a kinematically admissible displacement field for the given problem, b represents the body forces, and t the surface tractions. The solution u^* to

the elastostatics problem has minimum potential energy $\pi(u^*)$, where

$$\pi(u) = \frac{1}{2}\int_V \sigma : \varepsilon \, \mathrm{d}V - \int_V b \cdot u \, \mathrm{d}V - \int_S t \cdot u \, \mathrm{d}S. \tag{3}$$

The first term measures the internal strain energy induced in the body volume V; the latter two terms represent the external work performed by the body and surface loads, where the tractions are distributed over surface area S.

Note that the strain energy compares with the first-order quadratic stabilizers used in standard regularization [11–13]. Moreover, the elastic constants are related to the regularization parameter: by varying their values, we can modulate the stiffness of the body and thus the degree of smoothness in the deformations. It is evident that a range of regularizing behaviors may be effected by varying the particular idealized continuum on which the image volumes are modeled [14,15].

In elastic matching, the external forces are specified as the gradient of a similarity potential that is highly nonlinear in u. Iterative methods must therefore be used to solve for the equilibrium configuration. The basic strategy is to replace the nonlinear similarity term by a linearized approximation [1] and find the minimum $u^\star$ of the resultant quadratic expression for the total potential energy, as exemplified by (3). A new approximation valid near $u^\star$ is next developed, and the process is repeated until a stable estimate is reached. By formulating the statics problem using variational principles, the optimization at each iteration can be efficiently performed with the finite element method.

2.2 Finite Element Analysis

To solve for u^* in (3), only displacements of the form $u_h(x) = \alpha_i \phi_i(x)$ are considered from the set of kinematically admissible displacement fields. The use of a finite number n of trial functions ϕ_i to discretize the variational problem characterizes the Rayleigh-Ritz method [16]. The next step is to substitute u_h into (3). The new potential expression is a function only of α_i and its minimization,

$$\frac{\partial \pi}{\partial \alpha_i} = 0, \quad i = 1, \ldots, n,$$

yields a system of n equations for the n unknowns. Instead of solving the field equations at a finite set of points, the Rayleigh-Ritz method returns a continuous solution based on a finite set of functions.

[1] In practice, we use—as Broit did—a quadratic approximation to the nonlinear similarity potential, for which the total potential energy remains quadratic.

In the finite element method [17], the ϕ_i are defined piecewise according to subdivisions of the problem domain (called finite elements): in addition, they assume simple forms, such as low-degree polynomials, over any element; must be sufficiently smooth to ensure the relevant integrands are computable; and are designed to satisfy the relation $u_h(x_i) = \alpha_i$. The unknowns α_i are thus associated with discrete nodal points x_i, at which the elements are connected.

A primary feature of the finite element basis functions is that the unknowns at points within any element are interpolated using only the nodal values of that element:

$$u = N u^e, \tag{4}$$

where u^e is the set of nodal displacements of the element and the entries of N correspond to component parts of those basis functions with non-zero extent over the element. We can therefore write $\varepsilon = Bu^e$ within a given element, where B is composed of derivatives of N_i. The expression (2) for the stress tensor becomes, in matrix form, $\sigma = D\varepsilon = DBu^e$, over element e; D is the matrix of elastic coefficients. The total potential energy can be rewritten in terms of its elemental contributions and its minimization then calculated on an element-by-element basis:

$$\begin{aligned}\pi &= \sum_e \pi^e \\ &= \sum_e \left\{ \frac{1}{2} \int_{V^e} [u^e]^{\mathrm{T}} [B]^{\mathrm{T}} DBu^e \, dV - \int_{V^e} [u^e]^{\mathrm{T}} [N]^{\mathrm{T}} b \, dV - \int_{S^e} [u^e]^{\mathrm{T}} [N]^{\mathrm{T}} t \, dS \right\};\end{aligned}$$

$$\begin{aligned}\frac{\partial \pi^e}{\partial u^e} &= \frac{1}{2} \int_{V^e} [B]^{\mathrm{T}} DBu^e \, \mathrm{d}V - \int_{\mathrm{V^e}} [N]^{\mathrm{T}} b \, \mathrm{d}V - \int_{\mathrm{S^e}} [N]^{\mathrm{T}} t \, \mathrm{d}S \\ &= K^e u^e - f^e,\end{aligned} \tag{5}$$

where the stiffness and load characteristics of the continuous element have been "lumped" at the nodal points through K^e and f^e, respectively. Note that the nodal loads in f^e are unlike the pointwise forces used in finite difference methods [16]. The decomposition into finite elements is motivated in part by the efficiency with which the element stiffness and load terms can be constructed and evaluated in practice. To solve for the nodal displacements, the expressions in (5) are assembled into a large sparse system of equations. Additional information about the finite element implementation can be found in [18].

For images with sparse features, like those of anatomy, the finite element method allows parsimonious representations of the spatial mappings to be developed, which in turn can substantially simplify the matching calculation [19].

3 Bayesian Generalization

Broit's elastic matching method innovated the physics-based approach to shape modeling that is now widely applied in image analysis [20,21]. Its appeal and success derive from the available theory and numerical techniques for the continuum mechanics. On the other hand, the analysis in terms of mechanical systems is not the most general. In this section, we reformulate the image matching problem within a decision-theoretic framework based on *Bayesian modeling* and demonstrate some of its important features, including the computation of reliability and interval estimates and the incorporation of empirical information about the spatial mappings into the analysis.

Our formulation follows the general framework for Bayesian image analysis laid out in [22–25], and shares some elements of Grenander's pattern theory, which pioneered the probabilistic interpretation of flexible templates such as Broit's deformable atlas [26–28]. To illustrate the basic ideas, we begin with a general definition of the image matching problem: find x and f such that

$$I_s(X) = f(I_t(x(X)), \tag{6}$$

where the spatial transformation x specifies for each point X in I_s its corresponding point $x(X)$ in I_t, and f models the compound effect of processes which modify the values of the original intensity signal [29].

In our application, mappings are sought between images that arise from distinct but topologically similar sources, so (6) does not strictly apply—the success of atlas matching therefore depends both on the degree to which two anatomies actually bear topological resemblance and on the level of localization accuracy required by the analysis. Nevertheless, it is possible to identify in the images corresponding features whose registration will yield the desired mappings. By definition, each particular form of correspondence has associated with it a metric for quantifying the degree of similarity between two feature instances [29].

As an example, the following is a simple instance of (6) that is commonly assumed for magnetic resonance imaging (MRI) scans, after the original images have been suitably preprocessed:

$$I_s(X) = I_t(x(X)) + \mathcal{N}(\mu, \sigma^2), \tag{7}$$

where the noise process is additive and ruled by the same known Gaussian distribution at each pixel. In classical maximum likelihood (ML) estimation,

(7) leads to the following *likelihood function*:[2]

$$p(I_s, I_t; u) \propto \exp - \left\{ \frac{1}{2\sigma^2} \sum (I_t(X) - I_s(x(X)) - \mu)^2 \right\},$$

where the summation is performed over the image domain; I_s and I_t are realizations of the random fields modeling the image pair; and the displacement field, $u(X) = x(X) - X$, between them is treated as the unknown vector parameter of the distribution. For the imaging situation specified by (7), the (exponent of the) likelihood function above defines the pertinent criteria for rating the similarity induced between the observed images by any value of u. In particular, the mapping u^* which maximizes $p(I_s, I_t; u)$ is the the ML estimate.

Since the estimate effectively represents a *measurement* of the true disparity between the image pair [25], an error may be associated with each disparity estimate by examining the shape of the likelihood function [30]. The quadratic approximations to the similarity potential in section 2.1 amount to replacing the corresponding likelihood function with a Gaussian model at each iterative step.

3.1 Prior Information

For atlas matching, the estimated mapping between the images should not only have large probability as measured by the likelihood but must at some level be homeomorphic so that topological information about one image, the atlas, is preserved when transferred to interpret the other. These requirements, however, insufficiently constrain the problem because the image features on which matching is based are usually sparse, as, for instance, in anatomic scans where the strongest features lie primarily along the tissue interfaces. In section 2.1, we saw that by viewing the images as physical continua in the elastic matching, the associated strain energy or deformation measure essentially favors smooth mappings and so acts to regularize the matching problem.

The use of such *prior* information is a fundamental element of Bayesian methodology, and is introduced into the analysis by modeling the unknowns as random variables—the particular displacement field that we are after is considered to be a realization of the random field u. The prior knowledge is then expressed as a probability distribution on the space of admissible displacement fields.

For example, the presumption in elastic matching that the displacements should satisfy the governing equations of a linear elastic body can be encoded

[2] The observations at the pixels are assumed independent.

through a Gibbs distribution by making the potential proportional to the internal strain energy of the system [22,26]. For atlas matching, in addition to adopting such general assumptions about the spatial mappings, the analysis can be performed using actual statistics of the anatomic variation observed in a population.

To illustrate, assume the displacement fields $u^{(i)}, i = 1, \ldots, k$, mapping an atlas into anatomic alignment with a group of k subjects are available. The variance of the displacement mappings,

$$\mathrm{var}(k) = 1/k[u^{(1)} \ldots u^{(k)}][[u^{(1)}]^{\mathrm{T}} \ldots [u^{(k)}]^{\mathrm{T}}]^{\mathrm{T}} - \bar{u}^{(k)}[\bar{u}^{(k)}]^{\mathrm{T}},$$

computed in the Cartesian space will generally be singular because the number of observations in most practical situations will be far smaller than the dimensionality of the displacement fields; the latter have to be very high in dimension in order to model the complex morphological variations that occur naturally in human anatomy.

In order to derive from the observed variance an information matrix on which our prior model can be based, the statistics must therefore be represented in a subspace of the original Cartesian space. In particular, the subspace should grow in dimension as additional observations are introduced, and its basis must permit recursive updating. To satisfy these requirements, an orthonormal basis spanning only the space of the observed displacements can be constructed using the Gram-Schmidt process as follows:

Initial observation $u^{(1)}$:

- $u^{(1)} = w_1^{(1)} v_1, \quad \text{where} \quad v_1 = \dfrac{u^{(1)}}{\|u^{(1)}\|}, \quad w_1^{(1)} = < u^{(1)}, v_1 > = \|u^{(1)}\|;$

After $k-1$ observations, the orthonormal basis consists of $\{v_1, \ldots, v_r\}$:

- The new observation $u^{(k)}$ is projected on the basis $\{v_1, \ldots, v_r\}$: $w_{\mathrm{proj}}^{(k)} = [w_1^{(k)} \ldots w_r^{(k)}]^{\mathrm{T}} = [{v_1}^{\mathrm{T}} \ldots {v_r}^{\mathrm{T}}]^{\mathrm{T}} u^{(k)}$;
- If the orthogonal component, $u_{\perp}^{(k)} = u^{(k)}$ - $u_{\mathrm{proj}}^{(k)}$, is the null vector or is negligibly small, the basis is unchanged;
- Otherwise, the basis is augmented with the vector $v_{r+1} = u_{\perp}^{(k)} / \|u_{\perp}^{(k)}\|$: $u^{(k)} = [v_1 \ldots v_{r+1}][w_{\mathrm{proj}}^{(k)} \; w_{r+1}^{(k)}]^{\mathrm{T}}, \quad \text{where} \quad w_{r+1}^{(k)} = \|u_{\perp}^{(k)}\|.$

The initial basis vector v_1 is just the normalized version of the first observation and, like the remaining members of the basis, takes the form of a displacement field. As the displacements are processed, their mean and variance are simultaneously computed in the new basis representation, following the standard recursive scheme.

Our motivation for developing the new basis of reduced dimension was to circumvent the semidefinite nature of the variance matrix that results when the original Cartesian space is used. Under the new basis, it is possible that

$\mathrm{var}_V^{(k)}$ remains positive semidefinite, but the singularity will be limited to only a single null eigenvalue. Consider the eigendecomposition:

$$\mathrm{var}_V^{(k)} = G^{(k)}\Lambda^{(k)}G^{(k)^{\perp 1}},$$

where $G^{(k)} = [g_1^{(k)} \dots g_r^{(k)}]$ are the eigenvectors of $\mathrm{var}_V^{(k)}$ expressed in the orthonormal basis $V = \{v_1, \dots, v_r\}$, and $\Lambda^{(k)}$ denotes the diagonal matrix of eigenvalues. If there exists an null eigenvalue in $\Lambda^{(k)}$, it is removed; the corresponding eigenvector is similarly deleted from $G^{(k)}$. In Cartesian space, $\{g_1^{(k)} \dots g_r^{(k)}\}$ forms a new orthonormal basis of the subspace spanning the observed displacements:

$$[e_1^{(k)} \dots e_r^{(k)}] = [v_1 \dots v_r][g_1^{(k)} \dots g_r^{(k)}],$$

and its subset $\{e_1^{(k)} \dots e_{r\perp 1}^{(k)}\}$ spans the space of the displacement variance:

$$\mathrm{var}_{\mathcal{E}}^{(k)} = [e_1^{(k)} \dots e_{r\perp 1}^{(k)}]\Lambda^{(k)}[e_1^{(k)^{\mathrm{T}}} \dots e_{r\perp 1}^{(k)^{\mathrm{T}}}]^{\mathrm{T}}.$$

The mean $\bar{w}^{(k)}$, expressed in the basis $\{g_1^{(k)} \dots g_r^{(k)}\}$, becomes

$$\bar{z}^{(k)} = [g_1^{(k)^{\mathrm{T}}} \dots g_r^{(k)^{\mathrm{T}}}]^{\mathrm{T}}\bar{w}^{(k)},$$

from which its value in the Cartesian space can be derived as follows:

$$\bar{u}^{(k)} = [e_1^{(k)} \dots e_r^{(k)}]\bar{z}^{(k)}.$$

The mean $\bar{z}$ and variance Λ specify a Gaussian model for the observed displacements, with respect to the assumed atlas:

$$\begin{aligned} p(z) &\propto \exp -\left\{\sum_{i=1}^{r\perp 1} \frac{1}{2\lambda_i}(z_i - \bar{z}_i)^2\right\} \\ &\propto \exp -\left\{\frac{1}{2}\eta^{\mathrm{T}}\Lambda^{\perp 1}\eta\right\}, \end{aligned} \tag{8}$$

where $\eta = z - \bar{z}$, and $u = \bar{u} + [e_1 \dots e_{r\perp 1}]\eta$. This knowledge about anatomic variation, learned from a sample population, can be used to design a new matching procedure, where the search space is centered at the mean $\bar{u}$ of the previous observations and spans the principal modes $\{e_1, \dots, e_{r\perp 1}\}$ of the variance in those observations [31].

3.2 Likelihood Model

The Bayesian view of u as a random field alters the observation model accordingly: the likelihood is now formulated as the conditional probability $p(I_s, I_t|u)$

of observing the images given any particular value of the displacement mapping between them.

3.3 Posterior Analysis

Our prior beliefs $p(u)$ about the displacement mapping are revised when observations are available. From Bayes's law, the conditional density

$$p(u \mid I_s, I_t) \propto p(I_s, I_t \mid u)\, p(u)$$

describes the *posterior distribution* for u. Bayesian analysis proceeds from the posterior model and thus utilizes all the information that is available about the displacement field.

To perform matching, point estimates of u are inferred by minimizing an expected *loss* with respect to the posterior distribution [32]. The solutions obtained with Broit's elastic matching are, for instance, estimates of the mode for a particular posterior model and hence correspond to optimal actions under a zero-one loss function. To see this, compare (1) and (3) to the following log-posterior function:

$$\begin{aligned}\pi(u) &= -\log(p(u \mid I_s, I_t)) \\ &\propto -\log(p(u)) - \log(p(I_s, I_t \mid u)) \\ &= \frac{1}{2}\int_V \sigma : \varepsilon \, \mathrm{d}V - \int_V \text{similarity}(I_s(x), I_t(x+u))\, \mathrm{d}V,\end{aligned}$$

where the likelihood and prior models are given by the Gibbs distributions with potential proportional to the similarity norm and internal strain energy, respectively.

Another commonly cited point summary is the posterior mean, which optimizes the quadratic loss. To estimate its value, we sample from a series of normal approximations to the posterior distribution [33]. The specification of the loss function is a formal aspect of Bayesian decision theory and requires consideration of the cost incurred for making a particular decision given some value of the unknown parameters. Some of the practical implications of computing Bayes actions for the zero-one and quadratic loss functions in brain image matching are explored in [33].

In Bayesian analysis, a solution can be further characterized by its reliability so as to reveal the influence of the prior on the solution as well as the uncertainty in the observations [25]. The calculation of the second order statistics is made tractable in our implementation by considering only the covariance

between the components of the displacement vector at each point. The algorithm replaces the log-posterior function by a quadratic approximation valid near the given solution, and then computes the statistics of the corresponding Gaussian by Gibbs sampling [34].

The existence of the variance information admits another possibility: a confidence region or *Bayesian credible interval* may be developed about a point estimate. The result is an interval of possible displacement values for each point. In atlas-based localization, each structure can therefore assume a range of plausible shapes. An example is shown in figure 2, where we have additionally visualized the displacement variance on which the credible interval is based. Note that the underlying tissue interfaces are apparent in the variance plot, reflecting the fact that the uncertainty is reduced at the interfaces but is relatively large within those portions of the image that lack informative features. The same effect is evident in the narrowing of the interval estimate along the edges of the subcortical structures that line the ventricles.

4 Applications

From the outset our interest in image matching has been in its use to localize anatomy in cerebral scans of a subject by aligning an atlas with the subject's brain volume. In this section, we demonstrate how this approach enables an automated method for performing morphometry, and facilitates the analysis of functional activation studies made from groups of different individuals.

4.1 Atlas-based Morphometry

In addition to identifying a structure's location within the subject anatomy, the individualized atlas, as exemplified in figure 3, also yields direct knowledge about its size and shape. This information, when collected from a subject group, can be related to the anatomic variability among the individuals in that group.

Consider the images shown in figure 4 of the corpus callosum, that were extracted from midsagittal MRI sections of 8 male subjects. The corpus callosum is the band of fibers interconnecting cortical regions in one cerebral hemisphere with similarly placed regions in the opposite hemisphere. To account for differences in head placement within the scanner, each subject image was rigidly displaced so that three manually identified points in the image were brought into correspondence with their homologous counterparts in the atlas by least-squares optimization. Maximum *a posteriori* estimates of the spatial mappings

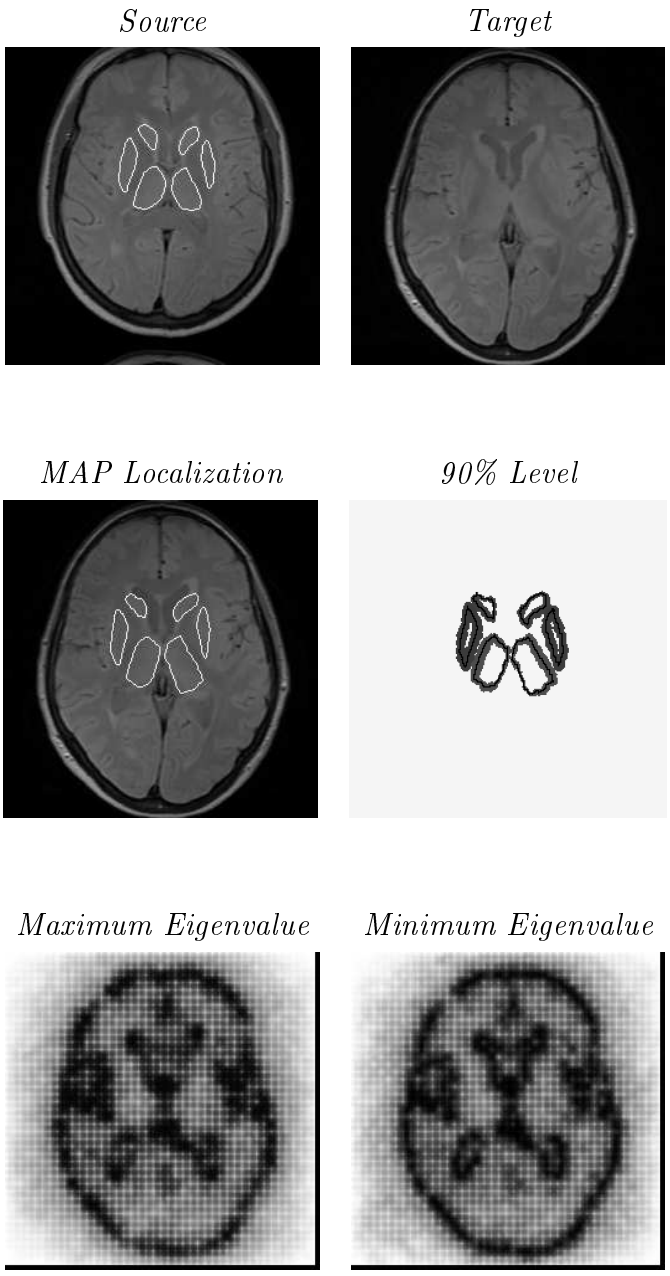

Fig. 2. Reliability and interval estimates for a maximum *a posteriori* (MAP) solution. Top row: Source image was matched to the target by MAP estimation; expert-defined contours of several subcortical structures for the source image were then superimposed on the target image using the estimated mapping. Middle row: 90% confidence regions (right) about the MAP localization (left). Bottom row: Reliability plots, where the intensity value at a point corresponds to the maximum (left) or minimum (right) eigenvalue of the variance in the estimated location for the point.

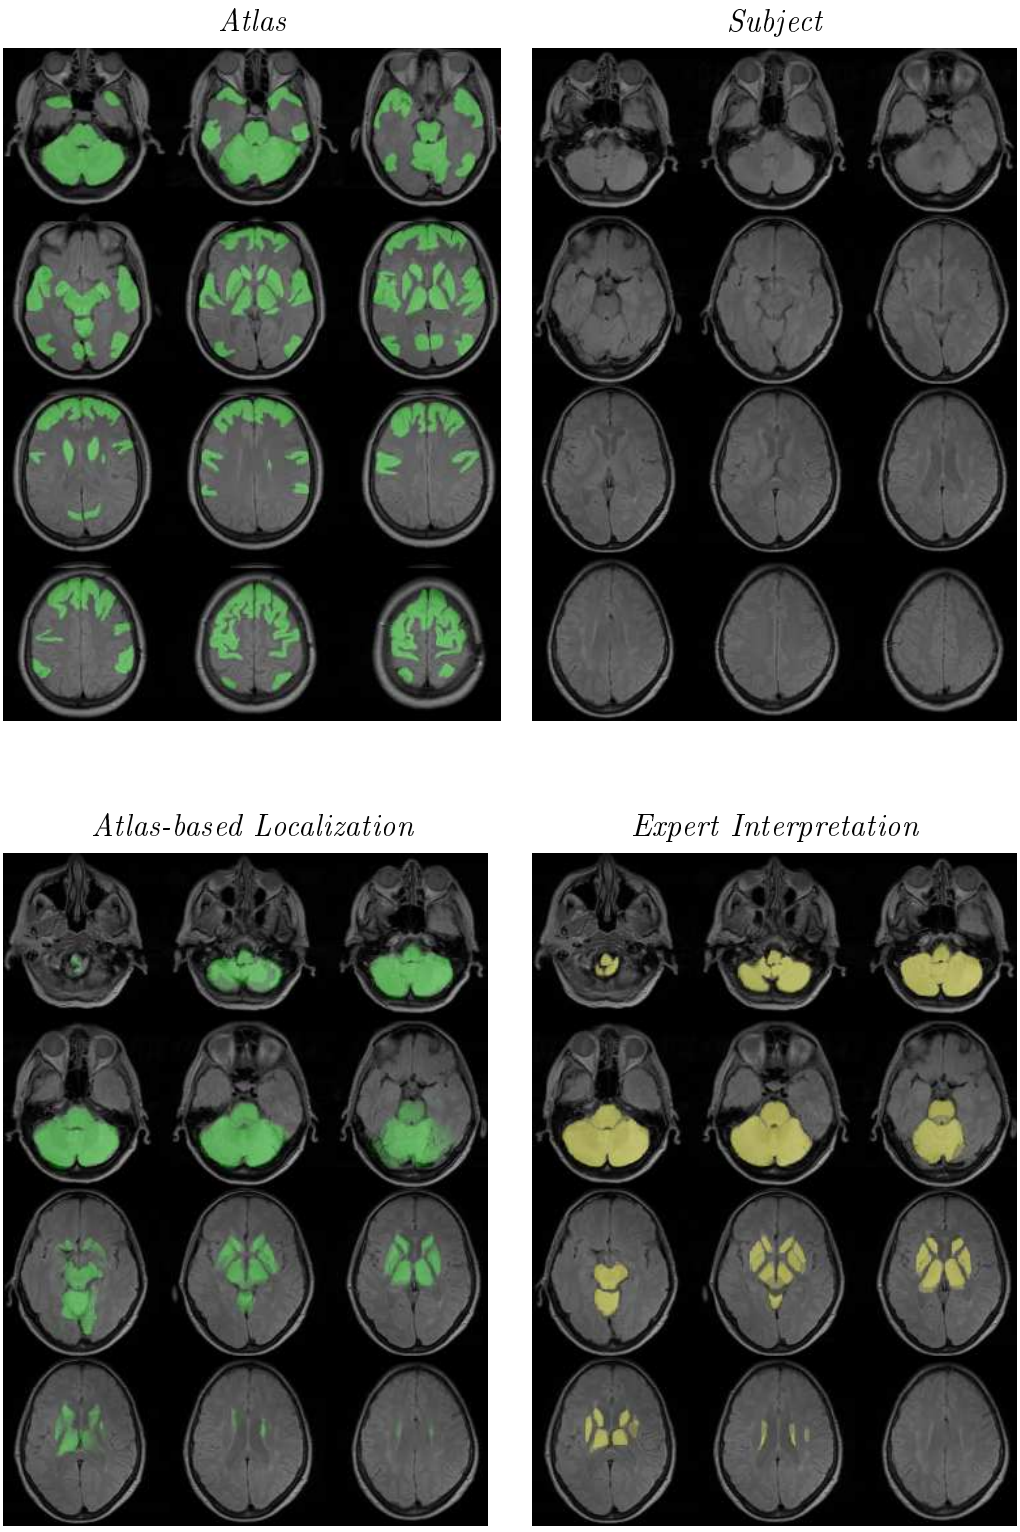

Fig. 3. MRI-based atlas localization of neuroanatomy. Anatomic labels, shown in green, encoded in the atlas (top left) are mapped to a subject's MRI study (top right), using the transformation obtained by deforming the underlying MRI volume of the atlas to match the appearance of the subject's MRI. The resultant anatomic localization (bottom left), superimposed in green over the subject, is compared with an expert's interpretation, shown in yellow (bottom right).

from the atlas to the subject group were then computed,[3] where the normalized cross-correlation between small neighborhoods centered about two given points was used to measure the similarity between them and the prior model was given by the Gibbs distribution with energy proportional to:

$$\int_V \sum_{i=1}^{2} \|u_i\|^2 \, \mathrm{d}V.$$

Figure 5 depicts the result for one of the subjects.

From the spatial mapping for a subject, one can derive the change in differential volume about any point in the atlas as the point is mapped into its corresponding position in the subject. The group average of these Jacobian determinants is shown in figure 6, where the areas of higher intensity value indicate dilatation and the darker regions correspond to compression of the atlas as it was warped into alignment with the subject callosa. As detailed in [35,36], the mean Jacobian at each point provides a useful measure with which to quantify the shape differences between two groups.

Toward a model that characterizes the callosal shapes of our subject population, the displacement fields were processed according to the procedure outlined in section 3.1. Figure 7 shows the average shape for the subject callosa in the group. The image was obtained by averaging together the individual mappings and then using the resultant mean warp to deform the atlas. The two most significant modes of shape variation within the group are also depicted in figure 7. To better illustrate the shape features modeled by each mode, the figure includes two sequences in which the mean callosal shape is deformed by scaled versions of the warps corresponding to the modes. The first principal component, with which 69% of the total variance is captured, measures the change in shape along the length of the corpus callosum, whereas the second principal component appears to contrast the size of the anterior and posterior portions of the structure.

[3] The atlas-based methodology emphasizes a fully volumetric approach *because* of the inherent difficulty of recognizing brain structures as they appear on image slices—this uncertainty is a consequence of the variety of shapes that a structure can assume depending on the position and orientation of the plane of section through its volume. The current analysis, on the other hand, utilizes two-dimensional data, which is special to the structure examined in this example. To eliminate the difficulty associated with anatomical localization on planar image sections, only the portion of the images containing the corpus callosum was used to guide the matching process; the external loads everywhere else were set to nil.

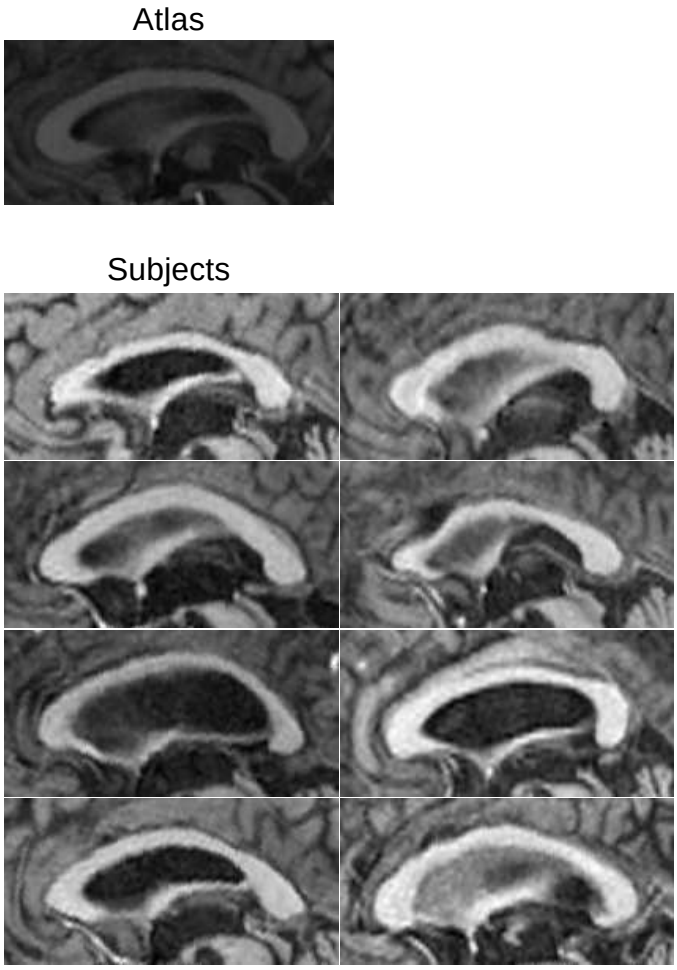

Fig. 4. Corpora callosa of 8 male subjects, along with the image of the corpus callosum used as the atlas or reference anatomy in the example.

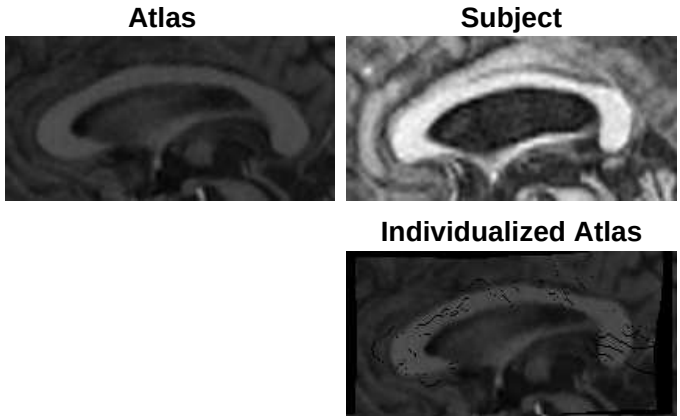

Fig. 5. Spatially aligned atlas for one of the subjects.

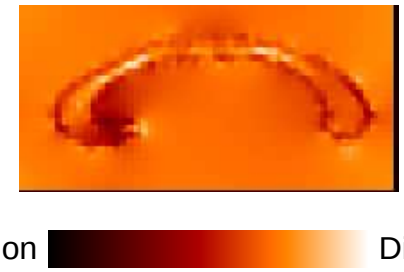

Fig. 6. Average change in differential volume at each point in the atlas as the atlas was aligned to the subject group.

4.2 Spatial Normalization for Functional Image Averaging

The capability of quantifying anatomic variations with respect to some reference space implies that these same variations may be removed in the analysis of group data. For imaging studies of brain activation, in which the task or stimulus induced activity are presumed to be localized to specific regions of the brain, the spatial precision with which brain structures can be aligned limits the highest resolution at which comparison can be made across a group of subjects to determine average functional responses or the variation of response within the group.

To accommodate the complex ways in which the neuroanatomy of normal individuals can vary, spatial transformations of very high dimension are essential. Consider the comparison shown in figure 8, where the structural image volumes [4] of 9 subjects from a functional MRI investigation of parahippocampal contribution in human topographical learning [37] were spatially normalized in two ways—one of the subjects was arbitrarily selected as the reference anatomy to which the remaining subjects were spatially aligned. In the first approach, the normalization was performed by estimating the 9-parameter affine transformations that best fit the overall size, location, and pose of the subject brains to those of the reference. In contrast, the second method applied a multi-resolution version of the matching algorithm [5] described in section 2.2, where each finite element estimate was allowed more than 40×10^3 degrees of freedom (distributed over a uniform hexahedral mesh spanning the image domain). The mappings inferred by each approach were then applied to the corresponding functional activation data, and the normalized activation images were finally averaged across the subjects.

[4] The original T1-weighted localizer images for each subject were reformatted for the current study into isotropic volumes (2 mm cubic voxels) with $x \times y \times z$ dimensions equal to $112 \times 96 \times 80$.

[5] A multi-resolution search for the MAP solution was performed using 3-D versions of the likelihood and prior models that were applied earlier for characterizing the callosal shape in a subject group. Additional details of the implementation can be found in [18,38].

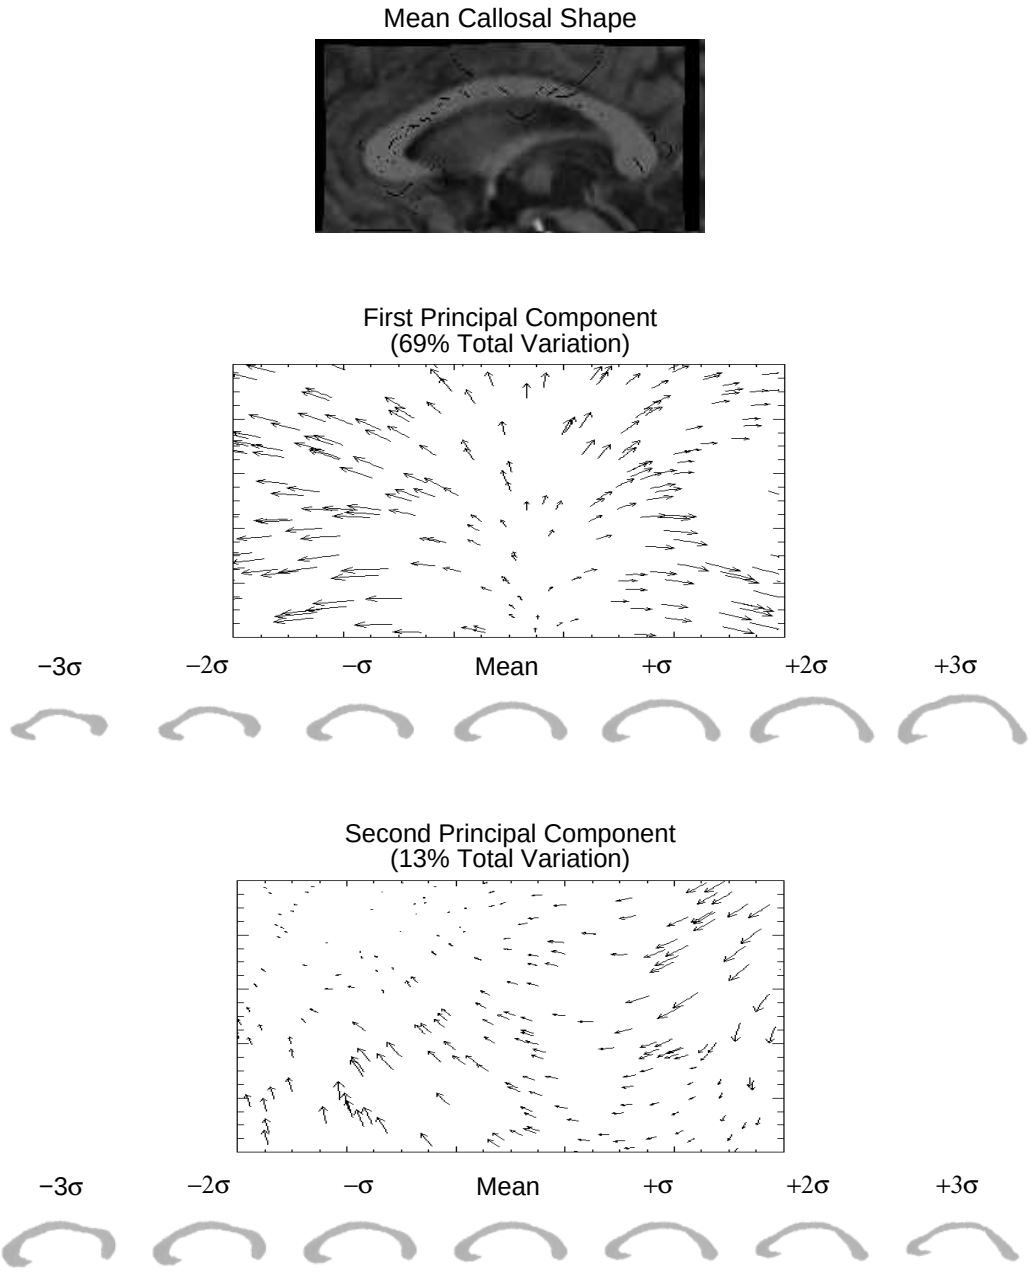

Fig. 7. Characterization of the callosal shape for the subject group shown in figure 4. To illustrate the shape features modeled by the first two principal components of the displacement variance about the mean deformation, the mean shape of the corpus callosum was warped according to the mapping shown for each component, where the amplitude of the component was scaled by the indicated number of standard deviations s.

The mean activation, shown in the figure overlaid on images of the average anatomy, was found to be higher in amplitude and qualitatively more spatially compact for the activation images aligned by high dimensional matching than for the affine normalized ones. Figure 8 also includes additional, higher resolution, sections through the mean anatomies: the superior normalization

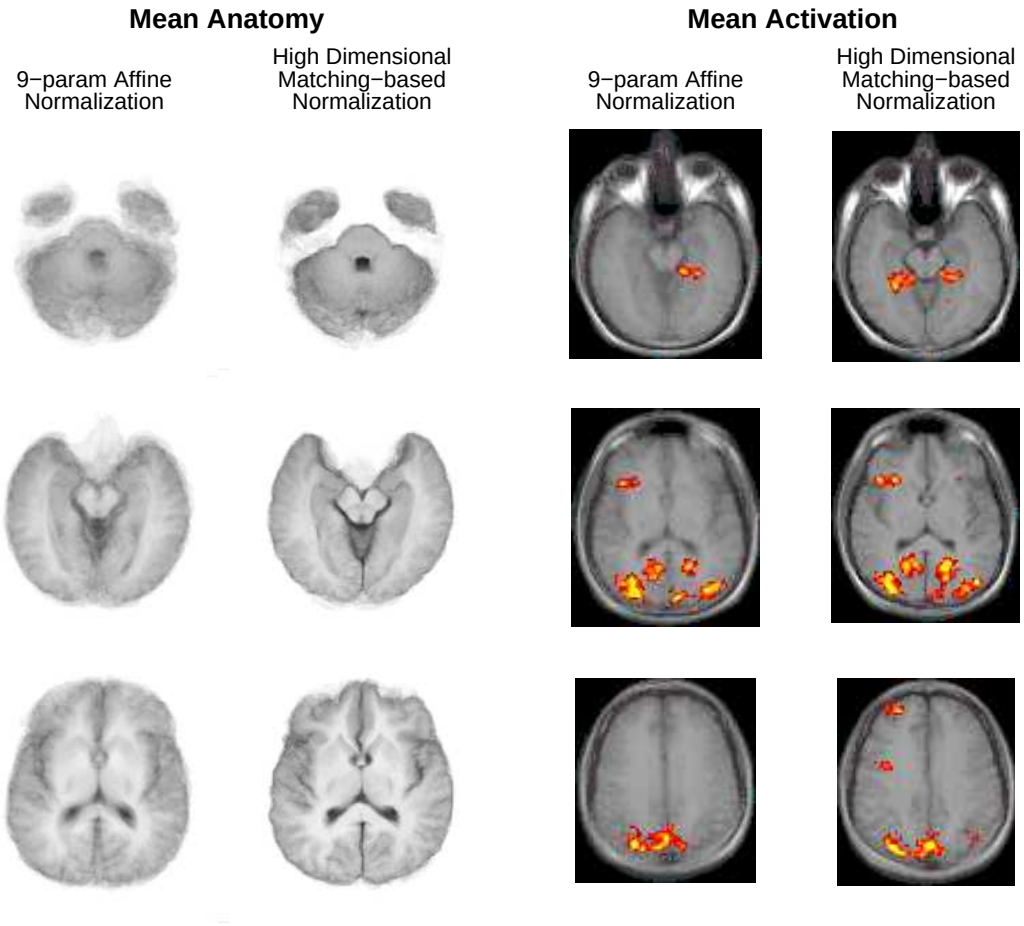

Fig. 8. Average activation and anatomy for a group of subjects as a function of the alignment method used for spatially normalizing the individual activation and structural images made from the group.

obtained with high dimensional matching is evidenced by the uniformly sharp appearance of its mean anatomic images. These results, detailed further in [38], highlight the necessity of transformations that are very high in dimension in order to extract the most information from group activation studies.

5 Summary

We examined the problem of image matching and its application to anatomic localization in volumetric images of the brain, introducing in the process an improved finite element implementation and more general Bayesian interpretation of the elastic matching method.

The finite element subdivision of a problem domain can be designed so that the precision of the computed solution varies spatially across the domain. For images of anatomy, where informative features are not only sparse but

distributed along interfaces of highly irregular shape, parsimonious representations of the spatial transformations are therefore possible, which in turn can substantially simplify the matching calculation.

We illustrated the role of elasticity in expressing our prior expectations about the nature of the anatomic variations among normal subjects, and constructed a posterior probability model for the displacement mappings, with which a variety of analyses can be conducted. In addition to the computation of reliability and interval estimates, a method was described for characterizing the observed morphological variation within a subject group in the form of prior distributions that can then be applied to guide the Bayesian analysis of future image studies.

Finally, a word is required about the case when the local variations in anatomy are very large in magnitude. To reduce the likelihood of false matches that arise from this situation, a standard approach is to solve the problem at different spatial scales, as in the multi-resolution version of elastic matching described in [39]: large scale displacements are first determined by matching the lower spatial frequencies in the images and these are then used to remove their confounding effect in the alignment of the higher frequency content. When even the multi-resolution search fails, additional external knowledge must be introduced into the problem formulation [40].

Acknowledgements

The work reported in this paper is part of a research effort that is marking its second decade at the University of Pennsylvania, during which many people have made important contributions. The author would like to make special mention of the following collaborators: Ruzena Bajcsy, David Haynor, Lionel Le Briquer, Alexei Machado, Pedro Peralta, Christian Barillot, Stane Kovačič, and Martin Reivich. The author is also grateful to Christos Davatzikos for sharing the corpus callosum data.

References

[1] C. Broit. *Optimal Registration of Deformed Images.* PhD thesis, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, 1981.

[2] R. Bajcsy and C. Broit. Matching of deformed images. In *6th Int. Conf. Pattern Recog.*, pages 351–353, 1982.

[3] F. L. Bookstein. *Morphometric tools for landmark data: geometry and biology.* Cambridge University Press, New York, 1992.

[4] M. I. Miller, G. E. Christensen, Y. Amit, and U. Grenander. Mathematical textbook of deformable neuroanatomies. *Proc. Natl. Acad. Sci.*, 90(24):11944–11948, 1993.

[5] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas: theory and rationale for its development. *Neuroimage*, 2:89–101, 1995.

[6] J. Declerck, G. Subsol, J.-P. Thirion, and N. Ayache. Automatic retrieval of anatomical structures in 3D medical images. In *Proc. Comput. Vision, Virtual Reality and Robotics in Medicine '95*, pages 153–162. Springer-Verlag, Heidelberg, 1995.

[7] D. L. Collins, A. C. Evans, C. Holmes, and T. M. Peters. Automatic 3D segmentation of neuro-anatomical structures from MRI. In Y. Bizais, C. Barillot, and R. Di Paola, editors, *Information Processing in Medical Imaging*, pages 139–152. Kluwer Academic, Dordrecht, 1995.

[8] P. M. Thompson, C. Schwartz, and A. W. Toga. High-resolution random mesh algorithms for creating a probabilistic 3D surface atlas of the human brain. *Neuroimage*, pages 19–34, 1996.

[9] C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Comput. Vision and Image Understanding*, in press.

[10] L. E. Malvern. *Introduction to the Mechanics of a Continuous Medium.* Prentice-Hall, Englewood Cliffs, 1969.

[11] A. N. Tikhonov and V. A. Arsenin. *Solutions of Ill-Posed Problems.* Winston, Washington, DC, 1977.

[12] T. Poggio, V. Torre, and C. Koch. Computational vision and regularization theory. *Nature*, 317(26):314–319, 1985.

[13] D. Terzopoulos. Regularization of inverse visual problems involving discontinuities. *IEEE Trans. Pattern Anal. Machine Intell.*, 8(4):413–424, 1986.

[14] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. Deformable templates using large deformation kinematics. *IEEE Trans. Image Process.*, 5(9), 1996.

[15] M. Bro-Nielsen and C. Gramkow. Fast fluid registration of medical images. In *Proc. Visualization in Biomedical Computing*, pages 267–276. Springer-Verlag, Hamburg, 1996.

[16] G. Strang. *Introduction to Applied Mathemetics.* Wellesley-Cambridge Press, Wellesley, 1986.

[17] K.-J. Bathe. *Finite Element Procedures in Engineering Analysis.* Prentice-Hall, Englewood Cliffs, 1982.

[18] J. C. Gee, D. R. Haynor, M. Reivich, and R. Bajcsy. Finite element approach to warping of brain images. In M. H. Loew, editor, *Medical Imaging 1994: Image Processing*, Bellingham, 1994. SPIE.

[19] J. C. Gee and D. R. Haynor. Numerical methods for brain warping. In A. W. Toga, editor, *Brain Warping*. Academic Press, San Diego, to appear.

[20] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Comput. Vision*, 1(4):321–331, 1988.

[21] D. Metaxas and D. Terzopoulos. Shape and nonrigid motion estimation through physics-based synthesis. *IEEE Trans. Pattern Anal. Machine Intell.*, 15(6):580–591, 1993.

[22] S. Geman and D. Geman. Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intell.*, 6:721–741, 1984.

[23] J. Marroquin, S. Mitter, and T. Poggio. Probabilistic solution of ill-posed problems in computational vision. *J. Amer. Stat. Assoc.*, 82(397):76–89, 1987.

[24] J. Besag. Towards Bayesian image analysis. *J. Appl. Stat.*, 16:395–407, 1989.

[25] R. Szeliski. *Bayesian Modeling of Uncertainty in Low-level Vision*. Kluwer Academic, Norwell, 1989.

[26] U. Grenander. Tutorial in Pattern Theory. Technical report, Division of Applied Mathematics, Brown University, 1983.

[27] U. Grenander, Y. Chow, and D. M. Keenan. *Hands: A Pattern Theoretic Study of Biological Shapes*. Springer-Verlag, New York, 1991.

[28] U. Grenander and M. I. Miller. Representations of knowledge in complex systems. *J. R. Statist. Soc.*, 56(4):549–603, 1994.

[29] L. Brown. A survey of image registration techniques. *ACM Comput. Surveys*, 24:325–376, 1992.

[30] P. Anandan. A computational framework and an algorithm for the measurement of visual motion. *Int. J. Comput. Vision*, 2:283–310, 1989.

[31] J. C. Gee and L. Le Briquer. An empirical model of brain shape. In G. Erickson, editor, *Maximum Entropy and Bayesian Methods*. Kluwer Academic, Dordrecht, 1997.

[32] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag, New York, 1985.

[33] J. C. Gee, L. Le Briquer, C. Barillot, D. R. Haynor, and R. Bajcsy. Bayesian approach to the brain image matching problem. In M. H. Loew, editor, *Medical Imaging 1995: Image Processing*, pages 145–156, Bellingham, 1995. SPIE.

[34] J. C. Gee, L. Le Briquer, C. Barillot, and D. R. Haynor. Probabilistic matching of brain images. In Y. Bizais, C. Barillot, and R. Di Paola, editors, *Information Processing in Medical Imaging*, pages 113–125. Kluwer Academic, Dordrecht, 1995.

[35] A. M. C. Machado and J. C. Gee. Atlas warping for brain morphometry. In K. M. Hanson, editor, *Medical Imaging 1998: Image Processing*, Bellingham, to appear. SPIE.

[36] C. Davatzikos, M. Vaillant, S. M. Resnick, J. L. Prince, S. Letovsky, and R. N. Bryan. A computerized approach for morphological analysis of the corpus callosum. *J. Comput. Assist. Tomogr.*, 20(1):88–97, 1996.

[37] G. K. Aguirre, J. A. Detre, D. C. Alsop, and M. D'Esposito. The parahippocampus subserves topographical learning in man. *Cereb. Cortex*, 6:823–829, 1996.

[38] J. C. Gee, D. C. Alsop, and G. K. Aguirre. Effect of spatial normalization on analysis of functional data. In K. M. Hanson, editor, *Medical Imaging 1997: Image Processing*, pages 550–560, Bellingham, 1997. SPIE.

[39] R. Bajcsy and S. Kovačič. Multiresolution elastic matching. *Comput. Vision, Graphics, Image Process.*, 46:1–21, 1989.

[40] S. C. Joshi, M. I. Miller, G. E. Christensen, A. Banerjee, T. Coogan, and U. Grenander. Hierarchical brain mapping via a generalized Dirichlet solution for mapping brain manifolds. In *Vision Geometry IV*, pages 278–289, Bellingham, 1995. SPIE.