ANIMAL: VALIDATION AND APPLICATIONS OF NONLINEAR REGISTRATION-BASED SEGMENTATION

D. L. COLLINS* and A. C. EVANS
Montreal Neurological Institute, McGill University
McConnell Brain Imaging Centre, 3801 University St.
Montreal, Canada H3A 2B4

Magnetic resonance imaging (MRI) has become the modality of choice for neuroanatomical imaging. Quantitative analysis requires the accurate and reproducible labeling of all voxels in any given structure within the brain. Since manual labeling is prohibitively time-consuming and error-prone we have designed an automated procedure called ANIMAL (Automatic Nonlinear Image Matching and Anatomical Labeling) to objectively segment gross anatomical structures from 3D MRIs of normal brains. The procedure is based on nonlinear registration with a previously labeled target brain, followed by numerical inverse transformation of the labels to the native MRI space. Besides segmentation, ANIMAL has been applied to non-rigid registration and to the analysis of morphometric variability.

In this paper, the nonlinear registration approach is validated on five test volumes, produced with simulated deformations. Experiments show that the ANIMAL recovers 64% of the nonlinear residual variability remaining after linear registration. Segmentations of the same test data are presented as well.

The paper concludes with two applications of ANIMAL using real data. In the first, one MRI volume is nonlinearly matched to a second and is automatically segmented using labels, predefined on the second MRI volume. The automatic segmentation compares well with manual labeling of the same structures. In the second application, ANIMAL is applied to seventeen MRI data sets, and a 3D map of anatomical variability estimates is produced. The automatic variability estimates correlate well ($r = 0.867$, $p = 0.01$) with manual estimates of inter-subject variability.

Keywords: Nonlinear deformation, warping, segmentation, variability analysis, human brain, MRI.

1. INTRODUCTION

Our interest in cerebral morphological variability stems from the work done in functional brain mapping at our institute.

In PET-based (positron emission tomography) cognitive activation paradigms, cerebral functional activity can be measured using a so-called subtraction paradigm. Here, two imaging experiments are completed; one during a resting state, and the second during some task requiring cognitive activity. The task can be as simple as tapping the right index finger or as complex as deciding on the ending consonant of monosyllabic words. The first image is subtracted from the second and the result indicates the regions activated by the particular task.

For some experiments, the signal difference between the active and the resting states is very subtle and averaging must be used to increase the signal to noise

*Corresponding author. E-mail: louis@bic.mni.mcgill.ca

International Journal of Pattern Recognition and Artificial Intelligence Vol. 11 No. 8 (1997) 1271–1294

ratio. Dose limitations prevent repeating PET studies more than a small number of times in a single subject, so the averages must be computed using multiple subjects.[1] Spatial registration between subjects is required for averaging, in order to align corresponding voxels, and is often achieved by mapping each brain volume into a standardized, or stereotaxic, coordinate space.[2] The use of some variant of the Talairach space[3] is almost universal, since it is based on easily identified anatomical landmarks, and accounts for orientation, position and scale.

While this procedure has been successfully applied in many studies, its limitations are now becoming apparent. Even after linear registration, there remains a considerable amount of anatomical variability unaccounted for. This residual misregistration is most notable for cortical structures (those on the periphery of the brain), which are also of interest in the analysis for cognitive processing. This misregistration is evident in the form of structure blurring when averaging images.

For example, the availability of a large MRI database of normal volunteers, acquired as part of an on-going brain mapping programme at the Montreal Neurological Institute (MNI), has lead to the construction of a 3-D probabilistic atlas of young, normal gross neuroanatomy, defined within stereotaxic space.[4,5] The average stereotaxic MRI volume shown in Fig. 1 was derived from 305 young normal individuals (239 males; 66 females; mean age 23.4 $\pm$ 4.1) after stereotaxic transformation of each MRI volume and intensity normalization. The blurring evident in Fig. 1 is a qualitative indicator of the local anatomical variability.

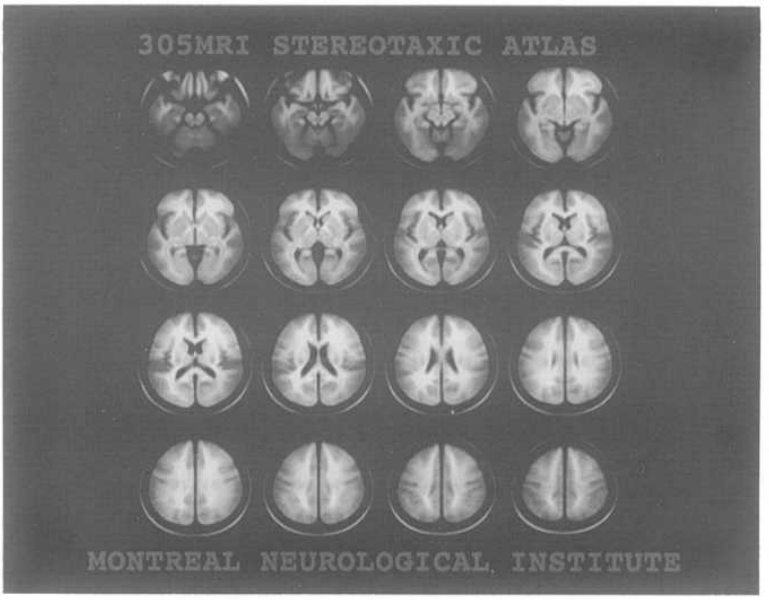

Fig. 1. Mean MRI dataset drawn from 305 young normal volunteers. The dataset is used as an anatomical template for locating functional activation data in Talairach space while respecting the known anatomical variability among individuals following linear stereotaxic transformation. It provides a visual impression of local anatomical variability and an indication of how well a particular functional measurement can be localized.

Unfortunately, the average stereotaxic MRI volume is insufficient for quantitative estimation of variability. For this purpose, the MRI intensity for each voxel in each MRI volume must be replaced by an anatomical label, e.g. caudate, precentral gyrus or calcarine sulcus, and a probability assigned for each voxel having a particular label. This requires a precise segmentation of each MRI volume into component structures, features and tissue types. However, manual labeling, in addition to being prohibitively time-consuming, introduces intra- and inter-observer inconsistencies in segmentation which can become comparable in magnitude to the true anatomical differences which are to be measured, confounding the overall goal. Our goal is to develop a method for completely automatic and accurate brain image segmentation at the regional level in order to remove these sources of error.

While averaging of labeled structures yields a probabilistic estimate for anatomical variability on a gross structural level, it does not give a measure for individual coordinates in terms of distance from some mean position. In the work completed by Sortie *et al.*,[6] 34 landmarks were identified on 17 brain volumes by 5 anatomists. Analysis of variability permitted the estimation of an inter-subject variability measure for each point, separated from intra-observer and inter-observer variabilities. The only difficulty with this work is that the measures of landmark variabilities are valid only for the landmarks chosen. Our goal is to develop a method that will permit the estimation of a dense field of variability estimates.

In this paper we describe the evolution of our work towards the two goals identified above, namely (1) automatic segmentation and (2) dense estimates of anatomical variability.

We have developed a nonlinear registration procedure known as ANIMAL (Automated Nonlinear Image Matching and Anatomical Labeling). In order to register a single subject to a given target, ANIMAL builds a 3-D nonlinear deformation field by sequentially stepping through the target volume in a 3-D grid pattern, estimating the displacement vector required to achieve local registration of the blurred gradient magnitude data extracted from each volume. The algorithm is applied iteratively in a multi-scale hierarchy, so that image blurring and grid size are reduced after each iteration, thus refining the fit. The recovered deformation field can then be used to resample the subject in the target space, or vice-versa, for direct voxel-to-voxel comparisons. Automatic gross structure segmentation is achieved by mapping labels, previously identified on the target model volume, through the deformation field, onto the subject's data volume. Since the anatomy is registered, the labels are essentially *registered* onto their corresponding structures to achieve segmentation. Estimates of variability for individual landmarks are derived from analysis of the recovered deformation.

This paper begins by describing similar work in the field of nonlinear registration, and follows by a detailed description of the ANIMAL procedure and the registration strategy. Experiments on automatic segmentation and analysis of anatomical variability are presented and validated using manual estimates. This paper ends with a discussion of the technique and presents directions for future work.

2. PREVIOUS WORK

Nonlinear registration is based on the assumption that there exists a topological equivalence between any two brain volumes. It is assumed that a continuous spatial transformation exists that can bring the two data sets into registration. (In the descriptions below, we will use the terms *source* and *target* to describe the first volume that is deformed onto the second, respectively.)

A number of groups have addressed this problem, and all the procedures developed involve (explicitly or implicitly) the identification of features, a measure of similarity based on these features and a method to define the nonlinear spatial transformation function. Here, some of the existing methods are grouped by the type of features used in the matching process: 0-D (points), 1-D (lines or curves), 2-D (surfaces) or 3-D (volumetric).

2.1. 0-D — Point Based Registration

Point-based registration requires the identification of homologous landmarks. Since the landmarks need not be organized on a regular grid, an interpolant for scattered data is required to define a continuous 3-D transformation function. The so called thin-plate-spline has been used in 2-D,[7] and 3-D,[8] to interpolate between manually identified landmarks to register brain volumes.

2.2. 1-D — Curve Based Registration

Curves have been used to increase the constraints used in the matching process. Ge *et al.*[9] subsampled manually extracted sulci and used the points to drive a thin-plate spline interpolant. Luo[10] also used manually identified sulci, but used a hierarchical force-based deformation approach to estimate the nonlinear transformation that best matches the sulci. Automatically extracted curves representing crest-lines of the ventricles are matched using an iterative closest point method by Subsol.[11]

2.3. 2-D — Surface Based Registration

Surface-based registration procedures require the extraction of homologous surfaces between the two volumes. The techniques differ in the way the surface is extracted and the way it is used to define the nonlinear transformation.

Downs *et al.*[12,13] define a pseudo-convex hull around the cortical surface. After linear alignment of the two data sets, radial interpolation is used to register the convex hull of the source volume on the target. While this type of method corrects for global shape differences between brains, it does not necessarily align gyri and sulci on the cortical surface.

Thompson *et al.*[14] have addressed this problem, where connected systems of parametric meshes are constructed within each brain to model major functional and lobar boundaries. Afterwards, the surface elements are warped onto the model brain using fluid dynamics to model the deformation process and extracted sulci are used as constraints to improve alignment of cortical structures.

Sandor *et al.*[15] also used cortical constraints to ensure registration of major cortical sulci. Morphological operators were used to identify sulci and gyri on an extracted cortical surface. Correspondence was achieved using a distance transform that was used to match nearest sulci together.

2.4. 3-D — Volume Based Registration

The techniques described above can be considered to be "label-based". These methods begin with explicit correspondence, achieved by the identification of homologous spatial structures. The 3-D nonlinear transformation is then computed that best superposes the extracted features. One problem with these methods resides in the choice of which features to use, how many to use and the subjective error when manually identifying equivalent landmarks in different brains. In previous work,[6] we have shown this uncertainty is on the order of 5–10 mm for point-based landmarks, similar to the extent of true anatomic variability that we wish to correct by using nonlinear registration. Another problem with these techniques is the choice of interpolant, since the deformation is estimated only at limited positions throughout the volume, the interpolant selected greatly affects the registration of the structures falling in between those used to calculate the transformation.

In the techniques described below, explicit correspondence is not available. It is derived by maximizing some indirect measure of similarity between the two volumes by optimization of the parameters driving the transformation. However, since data from the entire volume is used, the distance between the extracted features is much less than in the label-based methods above (usually on the order of the voxel spacing) and the choice of the interpolant becomes less important.

One of the first methods described in the literature is due to Broit *et al.*[16] In their program, registration between a 2-D computed tomography (CT) image and a corresponding atlas slice was achieved by physically modeling the atlas as a continuous elastic solid and deforming it to match the CT data. This technique has been used in 2-D to match CT images to a predefined atlas image,[17] and extended to 3-D.[18,19] Since uncertain quantities are involved in the matching process, Gee *et al.*[20,21] have introduced a probabilistic approach, where the Bayesian estimate of the transformation represents an optimal interpretation of the posterior model.

The group at the University of Washington also used a probabilistic formulation with physically based models to constrain the registration problem.[22] This work was extended to 3-D and included a Karhunen-Loeve model for linear elastic deformations.[23] Since the quadratic-based regularization models that were used are valid only for small deformations, Christensen *et al.* later replaced the linear elastic constraints with viscous fluid model in order to address the problem of large deformations.[24] The main difference between the two techniques is that stress restraining the motion relaxes over time in the fluid model, thus allowing large-magnitude deformations. From the same group, interesting work has been presented by Joshi *et al.*,[25] where point landmarks, curves, surfaces and volumes can be used together in a probabilistic setting of Bayesian inference problems.

As with our work presented here, the nonlinear registration technique developed by Friston *et al.* does not use a physical model to constrain the deformation.[26,27] Instead, the deformation is constrained to consist of a linear combination of smooth basis warps that are defined by discrete cosine transforms. The weighting factors (i.e. eigenvalues) are found that maximize a similarity criterion between source and target volumes. While the least-squares solution can be found quickly, the deformations recovered are limited by the predefined basis warps.

3. METHODS

3.1. ANIMAL: The Algorithm

The procedure described below is designed to recover the nonlinear transformation required to bring the two brain volumes into registration. Besides achieving nonlinear registration, the recovered transformation can be applied to both segmentation and dense estimates of anatomical variability.

3.1.1. The input

There are two data volumes used as input to the ANIMAL program. The first represents the *source* or subject's brain volume. The second is the *target* or model brain volume. For this paper, seventeen young normal subjects were scanned on a Philips Gyroscan ACS 1.5 Tesla superconducting magnet system. The data were acquired using T1-weighted 3-D spoiled gradient-echo acquisition with sagittal volume excitation (TR = 18, TE = 10, flip angle = 30°, 140–180 sagittal slices). Without loss of generality for the algorithm presented below, one of these subjects was selected to serve as the target. As part of an ongoing project at the MNI, each of the voxels in this volume was labeled at a gross anatomical level (e.g. superior frontal gyrus, thalamus, left lateral ventricle). Once labeled, this volume can serve as an atlas for the segmentation procedure described in Sec. 2.

3.1.2. The output

The goal of ANIMAL is to recover the nonlinear transformation required to register two brain volumes. The transformation is a spatial mapping function from $\mathbb{R}^3 \Rightarrow \mathbb{R}^3$. We assume that this function varies smoothly over the entire field and that it can be described as "locally translational", i.e. within a small neighborhood, the deformation field can be approximated by a translational flow field. In the implementation of the ANIMAL algorithm, the transformation is represented by a deformation field that is defined on a dense 3-D cubic lattice with a 3-D displacement vector stored for each node position in the lattice. In practice, three scalar volumes are stored: dx, dy and dz, representing the x, y and z-components of the 3-D displacement vectors. For a given arbitrary (x, y, z) position in the domain of the deformation function, the value of the corresponding 3-D displacement is given by interpolation in each component volume, yielding the three necessary values for the 3-D vector.

Since the deformation field is stored with finite sampling, only a band-limited representation of the continuous nonlinear spatial mapping function is possible. The scale of the deformation field depends on the scale of the estimation process. This is addressed in the next section.

3.1.3. Spatial frequency decomposition

It is natural to use large scale structures to establish large, low frequency deformations and conversely, small scale structures should be used to determine small, high frequency deformations. The deformations estimated from large scale structures can be considered to be an average of all the displacements affecting its smaller sub-components.

Intuitively, the inherent averaging represents a smoothed deformation field, and ties the scale of the data to the scale of the estimated deformation field. Therefore, we desire all computations to be scale dependent and this naturally leads to the use of scale-space methods.

We have selected a multi-resolution strategy (described in Sec. 3.1.7) where the deformation function is recovered hierarchically at a number of different steps, with each successive step refining the estimation of the previous one. Each step attempts to recover the deformation function for a given scale corresponding to the resolution of the band-limited deformation function to be recovered. If the FWHM of the current scale step is used to measure resolution, then the voxel spacing of the deformation field lattice must be no greater than FWHM/2 to recover the function without aliasing, i.e. the usual Nyquist sampling limit.

Note that at a scale of FWHM= 4 mm, there is a displacement vector every 2 mm. For a typical head size of $160 \times 150 \times 180$ mm, approximately 500000 vectors must be estimated (implying on the order of 1.5×10^6 degrees of freedom). At a scale of FWHM= 16 mm, less than 10000 vectors are required, so that a rough estimation of the deformation field can be computed in a fraction of the time (on the order of minutes) required for the final high resolution estimate.

3.1.4. The features

There are two constraints that must be satisfied by the features used in the estimation of the nonlinear deformation:

(1) Since all computations are scale dependent, the features used by ANIMAL should have a scalable property.
(2) In order to be independent of the original position, scale and orientation of source volume with respect to the target, the features must be geometrically invariant.

The original data sets do not satisfy these criteria, and so are not used directly by ANIMAL. We have chosen to use the blurred image intensity and image gradient magnitude so that the value of a feature corresponding to a particular anatomical landmark is the same regardless of its position or orientation within the image

volume. These features are calculated by convolution of the original data with zeroth and first order 3-D isotropic Gaussian derivatives. Convolution with such an operator maintains linearity, shift-invariance and rotational-invariance in the detection of features.[28,29] The Gaussian kernel also has a scalable property since it is dependent on the standard deviation, σ. In this paper, we use the full-width-half-max (FWHM= 2.35σ) of this kernel as the parameter to measure the spatial scale.

This parameter must be chosen carefully in relation to the scale of the estimated deformation function. A very large value will introduce too much blurring, possibly removing important structural detail. Too small a value will extract a lot of structure and consequently increase the probability of local mis-matches. The size of the Gaussian blurring kernel applied to the volumetric data was chosen to be equal to the resolution of the deformation field estimated at the current scale step. Figure 2 shows the gradient magnitude feature at scales of 16, 8 and 4 mm.

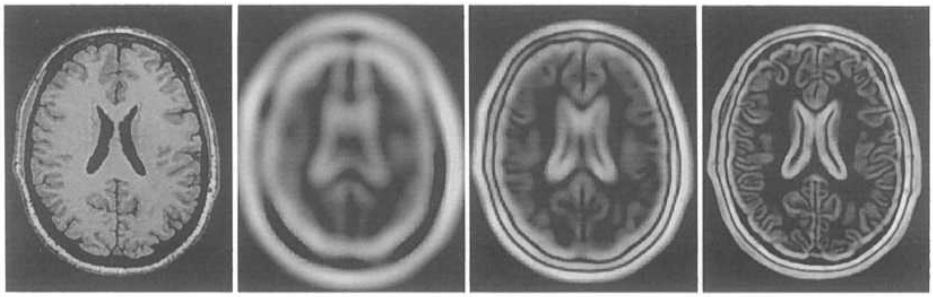

Fig. 2. Features used in registration. This figure shows a transverse slice at the level of the ventricles through the feature volumes used by ANIMAL. From left to right: The original data; the gradient magnitude feature at FWHM=16mm; FWHM= 8 mm; FWHM= 4 mm. One can see an increasing amount of detail as the scale is decreased.

3.1.5. The objective function

The objective function is used to evaluate the match between source and target volumes and is maximized by the optimization procedure. Ideally, the function would have a single maxima when the two volumes are in register, and be monotonically decreasing to be proportional to misregistration. In fact, the objective function has a complex shape with multiple local maxima. However, the use of the global multiresolution strategy described below minimizes the possibility of falling into a local maxima, as would be the case if the procedure were applied directly at the finest scale.

Here, the objective function has a classical formulation, that of a summation of similarity and cost terms:

$$S(\mathcal{S},\mathcal{T};\mathbf{N}) = \frac{1}{n}\sum_{\vec{x}\in\mathcal{L}}\left[R(\mathcal{S},\mathcal{T};\mathbf{N},\vec{x}) + C(\mathbf{N},\vec{x})\right]. \quad (1)$$

Where $\mathcal{S}$ and $\mathcal{T}$ are the source and target volumes, respectively; $\mathbf{N}$ is the nonlinear

transformation represented by the deformation field that maps points from $\mathcal{S}$ to $\mathcal{T}$; $R()$ is the local similarity measure and $C()$ is the cost function. The summation is evaluated over all nodes, $\vec{x}$, in the 3-D lattice, $\mathcal{L}$, of the deformation field, and is normalized by the number of nodes, n.

A number of functions are possible to evaluate the local similarity between the two data sets. We have chosen to use a normalized correlation statistic:

$$R(\mathcal{S},\mathcal{T};\mathbf{N},\vec{x}) = \frac{\sum_{v\in\mathcal{N}_{\vec{x}}} f(\mathcal{S},v) f(\mathcal{T},\mathbf{N}(v))}{(\sum_{v\in\mathcal{N}_{\vec{x}}} f^2(\mathcal{S},v))^{\frac{1}{2}} (\sum_{v\in\mathcal{N}_{\vec{x}}} f^2(\mathcal{T},\mathbf{N}(v)))^{\frac{1}{2}}}, \tag{2}$$

where $\mathcal{N}_{\vec{x}}$ is the local neighborhood of $\vec{x}$ with diameter = 1.5 FWHM, $f()$ is the volumetric interpolation function, and the summation is performed over all voxel elements $v \in \mathcal{N}_{\vec{x}}$. $R()$, and like-wise $S()$, take on a maximum value of 1.0 when the two volumes are in perfect registration.

The cost function introduces a penalty for large transformations, and is used simply to limit the maximum size of the estimated deformation vector:

$$C(\mathbf{N},\vec{x}) = \begin{cases} \dfrac{cd^{3/2}}{d_{\max}^3 - d^{3/2}} & \text{if } d^{3/2} \le d_{\max}^3 \\ \infty & \text{otherwise,} \end{cases} \tag{3}$$

where d is the length of the additional deformation vector required at node $\vec{x}$ and the constant $d_{\max}$ is chosen to be equal to the magnitude of the current FWHM. The constant $c = 0.2$ ensures that the function does not overly penalize small deformations.

3.1.6. Deformation field estimation

The deformation field is recovered in an iterative manner. Each iteration is separated in two steps: the first involves the actual estimation of the deformation vectors at each node by optimization of (1), and the second is a smoothing step to ensure a continuous deformation field.

Node estimation: In the implementation of ANIMAL, the transformation $\mathbf{N}$ is stored such that there is one deformation vector for each node $\vec{x}$. Since we define the nonlinear spatial registration transformation to be locally translational, the deformation at each node can be considered independent from its neighbors when taking into account the scale of the estimation step[a]. Therefore, (1) is maximized when each term in the summation is at a maximum. Hence, global optimization is divided into small mini-optimizations, where the goal is to find $\vec{d_i}$ that maximizes $R(\mathcal{S},\mathcal{T};(\mathbf{N}+\vec{d_i}),\vec{x}_i)$ for each $\vec{x}$ in $\mathcal{L}$ while minimizing the deformation cost $C(\mathbf{N},\vec{x})$.

The vector $\vec{d_i}$ is found using a three-dimensional Simplex optimization procedure, maximizing the correlation between $\mathcal{N}_{\vec{x}}$ in $\mathcal{S}$ and $\mathcal{N}_{\mathbf{N}(\vec{x})}$ in $\mathcal{T}$.

Smoothing constraints: Since a continuous deformation field is required, the estimation process must be constrained so that it cannot compress two distinct

[a]There is almost no overlap between the local neighborhoods for neighboring nodes.

points together or allow an overlap, nor can it induce a tearing of the field. This is achieved by tempering the deformation vector above with the average of the deformation vectors of the neighboring nodes. If $\mathbf{M}(\vec{x}_i)$ is the mean deformation vector calculated from the immediate neighborhood of $\vec{x}_i$, then the resulting deformation is given by:

$$\vec{d}_i' = \alpha \vec{d}_i + (1 - \alpha)\mathbf{M}(\vec{x}_i) \tag{4}$$

where $0 \leq \alpha \leq 1.0$. A small value of α ensures a very smooth deformation field at the expense of possibly missing some small local variations, while a large value of α gives more importance to the estimated deformation vector, with the risk of permitting local discontinuities to pass into the global deformation field. We have found that $\alpha = 0.5$ yields an acceptable balance between local matching and global smoothness.

Iterative refinement: At each scale step, the two step process is repeated iteratively on the partially deformed fields. Only a fraction of the estimated deformation vector for each node is applied and the estimation process is repeated until convergence. This approach dampens the tendency for artifactually large local shifts to introduce irreversible distortion in the deformation field. A fractional value of 0.6 is a compromise between speed and noise stability.

Since the optimization procedure terminates sooner when the initial conditions are close to the final result, and since calculating the fit at one scale requires less than one quarter the time taken for the fit at the previous scale, more iterations are allowed for the first two resolution steps: 15 iterations at 24 mm scale and 10 at 16 mm scale. Only 3–6 iterations are required at finer scales.

Node thinning: At each iteration, the deformation field is not calculated in regions with gradient magnitudes below a given threshold (set at 10% of the mean gradient magnitude), since no deformation can be reliably estimated in these regions. In this fashion, nodes that fall in the middle of a homogeneous region, away from any edge structures (e.g. the white matter of the centrum semiovale), are not used to define a local neighborhood and consequently, are not used to estimate a local deformation vector. The deformation for these nodes is interpolated from the neighboring lattice points in the smoothing step described above. This minimizes the influence of low signal-to-noise gradients and increases the speed of convergence.

It is important to note that the threshold is applied only to select or reject a particular node for estimation of the local deformation, and it is not used to eliminate any voxels from the calculation of local correlation for any selected node. Therefore, changing the value to 20%, for example, slightly thins most of the edges seen in the gradient magnitude image, leaving them behind so that they can be used to estimate the deformation. However, very low contrast edges may be lost, and thus the deformation in these regions must be interpolated from the neighboring nodes.

3.1.7. Coarse-to-fine multiresolution strategy

In order to speed up the estimation process and to avoid local minima in the objective function hyper-surface, the optimization procedure used to recover the best spatial transformations is accomplished in a hierarchical multi-scale fashion where the registration is performed at different spatial scales, starting with very blurred data and increasing detail at each step by using less blurred images, refining the registration at each stage.

The volume field-of-view dictates the absolute largest scale of the deformation and the smallest is limited by the resolution of the data. A factor of 2 is classically used to step through scale space, starting at the original pixel size and increasing by 2, 4, 8 and 16 times the original pixel size. Initial experiments showed that kernels with a FWHM= 32 mm blurred the data too much to be of any use, so the nonlinear registration procedure begins by using data blurred with a FWHM= 24 mm Gaussian kernel and lattice node spacing of 12 mm. After ANIMAL is applied at this scale, the resulting deformation field is resampled onto a 8 mm lattice, and the estimation continues using data at FWHM= 16 mm. Depending on the application, the process can stop after the 16 mm fit, or continue decreasing scale with data blurred at 8, 4 and 2 mm for higher resolution fits and consequently longer run times.

On average, when running on an SGI Indigo^2 (175Mz R4400 CPU, 115 SPECfp92), each iteration at each scale takes 2.4 minutes at FWHM= 24 mm with a 3D lattice containing approximately 7600 nodes, 3.6 minutes at FWHM= 16 mm (28600 nodes) and 12.2 minutes at FWHM= 8 mm (253000 nodes).

3.2. Segmentation

Manual anatomical segmentation is usually accomplished by comparing a figure from an anatomical textbook with the digital image to be segmented. The position of a particular border is refined by comparing it, and its neighboring structures, to the contours in the atlas. Hence, segmentation is essentially a registration problem.

A similar process is applied here. Since the source and target volumes are registered, the predefined labels of the target volume inherently segment the source volume as well. Explicit segmentation of the source volume in its native space is achieved by mapping the labels through the inverse transformation recovered by ANIMAL.

Implicit in this analysis is the assumption that the first data set is perfectly segmented and that the one-to-one mapping exists. Even though the latter is not strictly true, one brain may have a sulcal pattern that another does not, the assumption does not lead to catastrophic deformations of one structure onto another in practice. The brains are matched at successively finer scales until the assumption of anatomical correspondence breaks down.

The main advantage of this type of segmentation methodology is that it results in *atlas-independent segmentation* since structure identification is simply a by-product of nonlinear registration. The anatomical labels defined on the target volume are not used in any way to determine the match between data and model. Therefore, any atlas defined on the target MRI brain volume can be used for segmentation, thereby

allowing for the co-existence of multiple atlases, each of which is simultaneously mappable to the native MR image volume without the CPU-expensive recalculation of the nonlinear spatial transformation required for registration.

3.3. Variability Analysis

The analysis of anatomical variability requires a common coordinate system with respect to which the spatial variability of corresponding features from different individuals can be expressed. Since that positional variability will depend on the degrees of freedom (DOF) allowed for mapping the image data from its original, native space into stereotaxic space, it is essential that the transformation from the native to the common space be well-defined and reproducible. For simplicity and ease of use of interpretation, it is also desirable that the transformation has as few DOF as possible while accounting for most of the positional variability.

For this study, a Talairach-like brain-based coordinate system has been chosen as the standard,[3] so called stereotaxic, space and a simple 9-parameter linear transformation is used to map brain volumes into this space.[30] Any discussion of measurements, averages or variabilities must bear these two choices in mind.

It is important to note that while they are related, the problem of *quantifying* anatomical variability is different from the practical issue of *removing* anatomical variability. The goal of latter is to render all brains identical after transformation, and can be achieved (in part) by the application of the nonlinear spatial transformation recovered by ANIMAL, described earlier.

To quantify anatomical variability, it is necessary to identify homologous features between the source and target volumes and to measure the difference in position between them, within the stereotaxic space. This identification can be achieved automatically by ANIMAL, since the recovered deformation field essentially establishes correspondence between the homologous points. The field can be interpreted as the map of "positional differences" between individual source (after affine transformation) and target volumes. In other words, for every 3-D coordinate in the target space, the registration procedure yields a vector-valued estimate of the difference in position between the two data sets. Thus, the two goals of identification and difference measurement are met by ANIMAL.

This information derived from the deformation fields, when averaged over a large number of individual/target pairs, can yield estimates of normal anatomical variability. The standard deviation at each voxel position (over the deformation fields estimated for the 17 subjects described in Sec. 3.1.1) is computed separately for each of the x, y and z components. These values are combined to yield a single number for each voxel measuring inter-subject variability (ISV), equivalent to a 3-D FWHM measure: $\text{FWHM} = 2.35\sqrt{(\sigma_x^2 + \sigma_y^2 + \sigma_z^2)/3}$.

4. EXPERIMENTS AND RESULTS

Three experiments are presented to validate the ANIMAL algorithm. The first involves only simulated data to examine nonlinear registration and brain structure

segmentation in a completely controlled setting. In the second experiment, *real* MRI data is used to demonstrate the power of ANIMAL for both registration and segmentation in a real-world situation. Finally, a comparison between the manual and automatic estimates of anatomical variability is shown in the last experiment.

4.1. Simulations

In order to validate the nonlinear registration algorithm with real volumetric MRI data, a single data set was selected from the original group of 17. A random 3-D deformation was applied to the selected brain volume to create a warped data volume. ANIMAL was used to recover the applied transformation.

The random 3-D nonlinear transformation was defined in the following manner. An original set of twenty landmarks were selected throughout the brain volume. A deformed set of landmarks was produced by adding a random displacement to each landmark coordinate. Previous work in our laboratory has shown anatomical variability to be on the order of 4–7 mm.[6] A Gaussian random number generator with a standard deviation of 5 mm was used to produce each component of the displacement vector.

The two resulting point sets were then used to define a continuous 3-D thin-plate spline (TPS) transformation function. This interpolant was used to resample the original data and produce the spatially warped source data set for testing. The average deformation magnitude is on the order of 7.7 mm with a maximum of 19.7 mm.

Figure 3 shows transverse slices through the original and five warped test volumes. While these images demonstrate extreme deformations, they form a good test-base for ANIMAL.

4.1.1. Nonlinear Registration

The multiresolution ANIMAL registration strategy was applied to each test-original source/target pair, and the recovered transformations were used to resample the test volumes back into the original target space.

Figure 3 shows the resulting volumes after linear registration. While each head is approximately the same size and orientation, the transverse slices demonstrate that they are not at all similar in shape, nor are the brain structures sliced at the same level within each brain. Figure 4 shows the same brains after nonlinear registration. It is apparent that the deformation field recovered by ANIMAL has accounted for a great deal of the shape nonuniformity. The ventricles appear to have the same shape, the basal ganglia structures are sliced at the same level and much of the cortex is properly aligned. The average images in Fig. 5 were created by simply computing the mean intensity value over the five brains for each pixel position in the volume. Here it is clear that the nonlinear deformation field has been recovered over the five test cases.

Fig. 3. Data for experiment 1. This figure shows transverse slices through the level of the ventricles for the original (upper left) and the five warped test volumes, after resampling with the linear transformation recovered by ANIMAL. These volumes are perhaps deformed more than one would find on average in the normal population, however subjects representing an extreme of normal anatomical variability could be shaped like these examples. Note that while only 2-D images are shown in the figures, all calculations are computed in 3-D on volumetric data. Deformation procedures that are applied only on 2-D images from a volumetric data set can never recover the structure that is out of plane, regardless of the amount of 2-D deformation permitted.

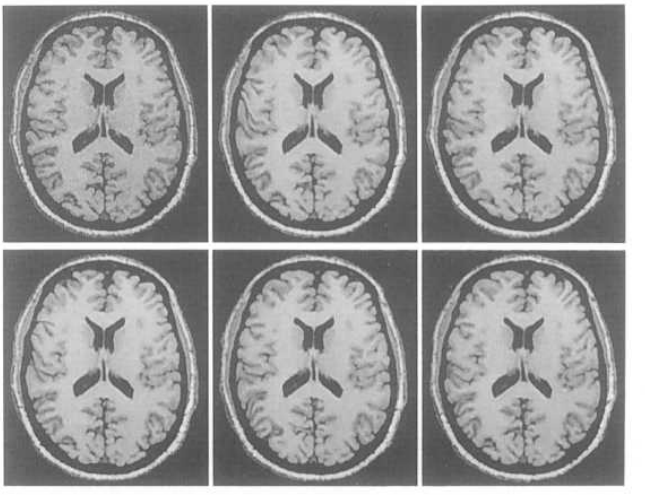

Fig. 4. Results of experiment 1. This figure shows the original (upper left) and the five warped volumes, after resampling with the nonlinear deformation field recovered by ANIMAL. One can see that not only the global head shape is accounted for, but internal structures such as the ventricles are well aligned.

In order to better judge the quality of the recovered transformation, the resampled test volumes were subtracted from the original volume and the resulting difference images are shown in Fig. 6. Here, one can see that ANIMAL does indeed

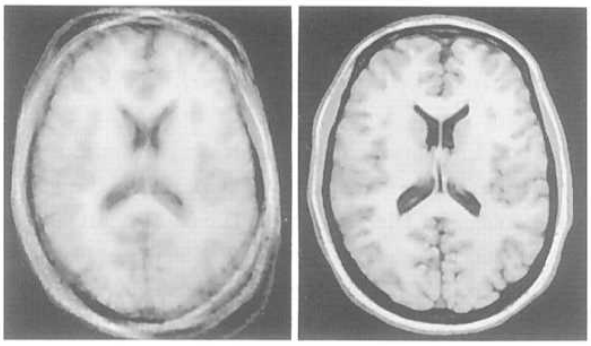

Fig. 5. Average images: linear vs nonlinear. This figure shows an intensity average of the resampled test volumes. On the left, the five volumes were resampled by the linear transformation and on the right, by the non-linear transformation recovered by ANIMAL. One can clearly see the advantage of the nonlinear registration when pixel alignment is required for image processing.

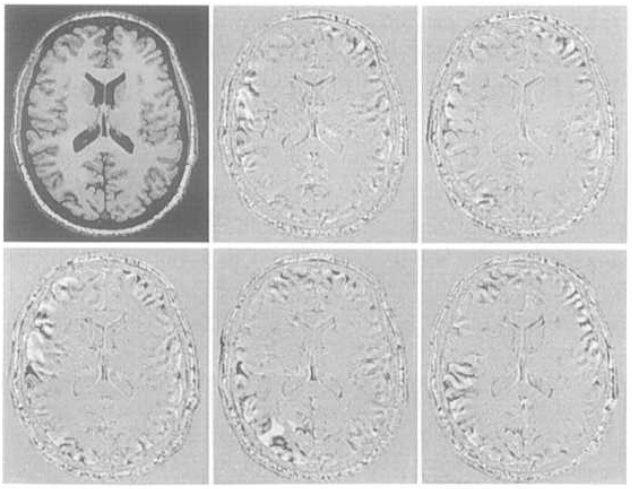

Fig. 6. Difference results for experiment 1. These difference images were computed by subtracting the resampled test volumes from the original volume (upper left). The pixel intensity scale of the difference images is six times greater than the original MR images. One can see that in general, the nonlinear registration is quite good however some regions of the cortex appear to be mis-aligned (indicated by very bright or very dark regions).

properly recover the transformation required to register source and target volumes. However, there are limited regions near the cortex that are misaligned.

Quantitative evaluation of ANIMAL is possible in this experiment since the spatial transformation between the source and target data sets is known. Indeed, the ability of ANIMAL to recover the applied transformation can be measured for specific landmark points within the volume. In order to have a brain-size independent measure, we have chosen to define a "recovery error" in the following way. A 3-D coordinate p, defined in the target space is first mapped through the transformation used to produce the warped volume, yielding p'. The resulting coordinate, p', is

then mapped back into the target space using the recovered transformation to give p''. If the recovered transformation were perfect, then there would be no difference between the positions of p and p''. The recover error is defined to be $r = |p - p''|$, where $|\cdot|$ is the l_2 norm.

In order to have a good idea of the registration throughout the brain volume, a 3-D lattice of points with 10 mm spacing was defined on the target brain (yielding 2010 nodes). On an average for the 5 volumes, the root mean squared (rms) recovery error evaluated over all nodes was 5.6 mm after linear registration. This error was reduced by 64%, to 2.0 mm, after nonlinear registration.

4.1.2. Segmentation

As described in Sec. 3.2, gross structure segmentation is achieved by mapping labels, previously defined on the target volume, through the inverse of the recovered nonlinear deformation, onto the source volume. In this way, segmentation can be used as another tool to validate the nonlinear registration process when compared to an expert's manual segmentation of the same structures.

A neuro-anatomically trained expert identified cortical grey matter structures (superior frontal gyrus, middle frontal gyrus, inferior frontal gyrus), basal ganglia structures (thalamus, caudate nucleus, putamen, lentiform nucleus) and cerebrospinal fluid (CSF) filled spaces (lateral ventricles) on the selected target volume. These structures were mapped through the inverse of the recovered transformation onto each of the five test volumes and the resulting segmented images are shown in Fig. 7. One can see that the structures in the test volumes are well delimited by the recovered transformation.

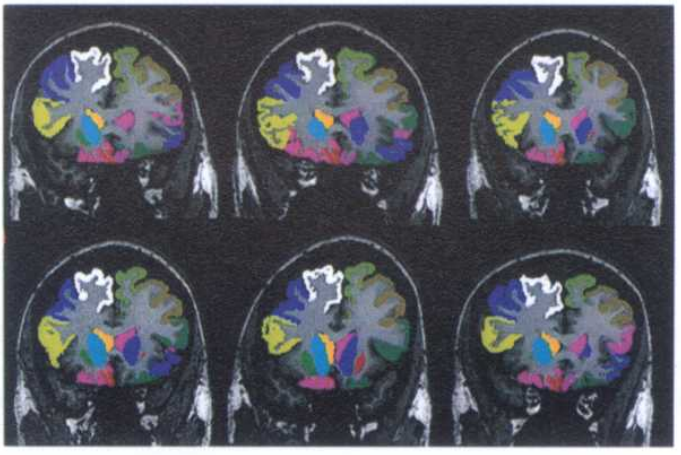

Fig. 7. Segmentation results for experiment 4.1.2. This figure shows the structure labels on the original (upper left) and the automatic segmentation on the five warped data volumes (in their native space). The segmentation results from mapping the original labels through the inverse of the recovered transformation. One can note that even the cortical structures are well segmented by the ANIMAL algorithm.

4.2. Real MRI Data

Since ANIMAL has been shown to work well on real MRI data from the same subject with simulated deformations to model other subjects, the next logical step is to investigate its application to register and segment real MRI data from different subjects. We have continued to use the previously described labeled brain as the target and a second MRI volume was selected from the 17 subjects described in Sec. 3.1.1 and used as the source volume. ANIMAL was applied to recover the source to target transformation.

The top row of Fig. 8 shows slices through the target, the source after resampling through the recovered linear transformation and the source after nonlinear transformation. One can see that ANIMAL does indeed not only recover the global shape of the head and brain, but more importantly aligns internal structures to enable voxel-to-voxel comparisons. In order to judge the quality of the registration, the bottom row of Fig. 8 shows the difference images between the target and the source volume, after linear and nonlinear resampling. These images confirm that the nonlinear registration has accounted for the anatomical variability between the source and target brains.

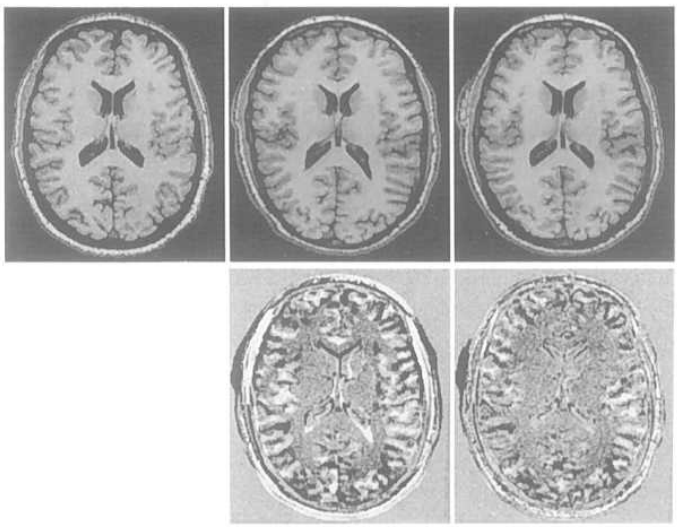

Fig. 8. Registration results for real data. Top row left to right: target MRI, source MRI after linear registration, source MRI after nonlinear registration. Bottom row: corresponding difference images. Here one can see that even when applied to real data, ANIMAL is able to account for brain shape and structure position.

Once again, the target volume labels were mapped through the inverse of the recovered transformation to achieve segmentation of the source volume in its original space. The structures defined on the target were also labeled by hand on the source by a neuroanatomical expert and were used to validate the automatic segmentations. Figure 9 shows the comparison between the manual and automatic segmentations.

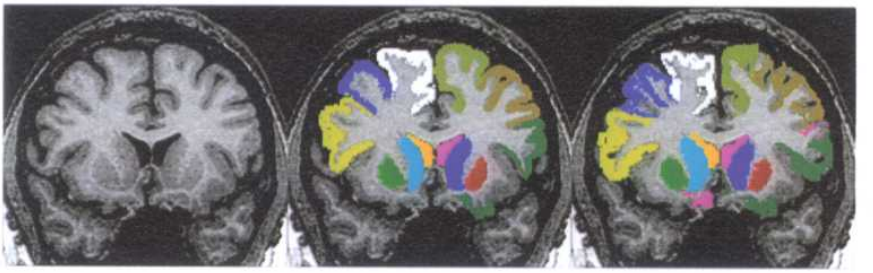

Fig. 9. Segmentation results for real data. Left: a coronal slice through the frontal cortex of the source volume; Middle: with manually identified structure labels; Right: with automatically identified structure labels.

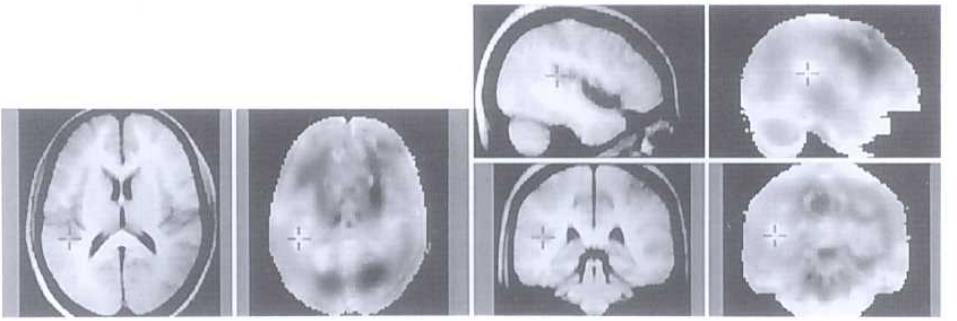

Fig. 10. Variability map. The images show the average intensity volume of the 17 subjects mapped into stereotaxic space nonlinearly along with the corresponding slices through the average variability map. The cross marker ($x = -44$ mm, $y = -37$ mm, $z = 14$ mm) is near the planum temporale, a region known to be variable, measured here to be 6.3 mm 3-D FWHM, and appearing more variable than the right.

4.3. Anatomical Variability

After automatic segmentation, the second goal set forth in the introduction was the automatic estimation of anatomical variability. The method described in Sec. 3.3 was applied to a group of 17 MRI brain volumes (3-D spoiled gradient-echo acquisition, TR = 18, TE = 10, flip angle = 30°, 1mm^3 voxels), registering them to a MRI-based stereotaxic atlas.[5] The 17 deformation fields were used to compute the anatomical variability map shown in Fig. 10. The regions of largest neuroanatomical variability were posterior poles of the lateral ventricles, the region near the fourth ventricle, the cingulate sulcus (slightly more on the left than the right), the inferior frontal lobe and the area just above the splenium of the corpus callosum. The anatomical variability map is not symmetric on the left and right sides. The left frontal lobe and the right parieto-occipital lobe appear to be more variable than their counterparts.

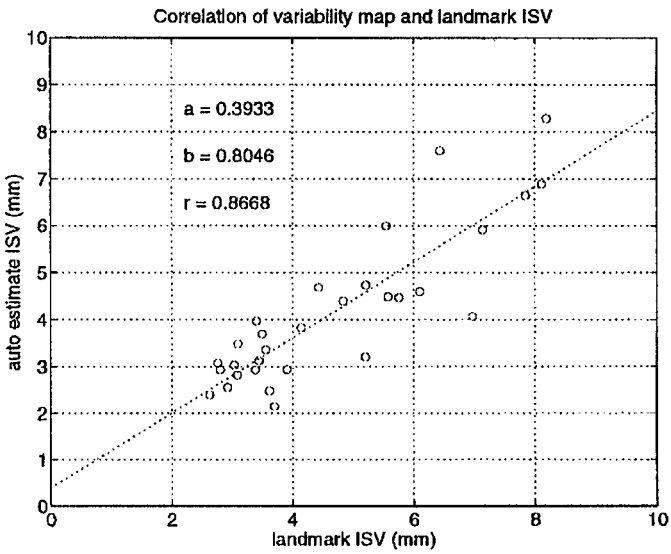

Fig. 11. Correlation between manual and automatic variability estimates. This graph shows the correlation between the automatic estimate of neuro-anatomical variability and the manual estimate based on ISV for each of the 34 landmarks. The regression coefficient is 0.867, and is significant at the 1% level.

For quantitative validation, the manual ISV estimates from a previous project[6] were compared with the automatic estimates derived here and plotted in Fig. 11. The average manually estimated ISV value was 3.9 mm compared to 4.2 mm here. The regression coefficient was 0.867, demonstrating good correlation between both methods at the 1% significance level.

5. DISCUSSION

The experiments presented in Sec. 4.1 indicate that ANIMAL is capable of recovering important deformations between the source and target volumes (at least up to the 20 mm tested), resulting in almost perfect registration after resampling through the recovered transformation (see Fig. 4). It is important to note that the structure in the difference images of Fig. 6 may be due to causes other than mis-registration. Trilinear resampling was used to map the test volumes back onto the target before the subtraction, so the very small (thin) edges seen in the difference images may be due to interpolation errors. Also, the ANIMAL procedure was stopped at the 8 mm scale (i.e. a deformation vector every 4 mm). Perhaps the procedure should be applied again at a finer scale to recover the smallest residual deformations not accounted for at the 8 mm scale.

The recovery error is dependent on the exact point-to-point correspondence throughout the entire volume. Since this value was non-null in the simulations, it may indicate a limitation of the assumption that local neighborhood correspondence can be used to recover the global transformation. Indeed, deformation vectors can

be estimated only where features were extracted from the volumetric data sets. In the implementation of ANIMAL, vectors are estimated only in regions where the gradient magnitude is greater than a small threshold. For homogeneous regions (i.e. where the gradient magnitude is very small), the deformations are interpolated from neighboring regions in the smoothing step. This problem affects not only the ANIMAL registration procedure but most registration procedures in the literature that are based on features extracted from the data.

Still, the recovery error does indeed quantify the difference between the applied (i.e. simulated) and the recovered nonlinear deformations. While this measure is important to validate the technique and is important in the measure of anatomical variability it may be perhaps too strict for other applications. For example, point-to-point correspondence is not strictly required for structure labeling. As long as equivalent labels are mapped together, segmentation is achieved as is shown in Fig. 7 where even the cortical structures are well segmented. These experiments show that registration can indeed be used to achieve segmentation.

The major difference between the simulations described above and the experiment using two different subjects for the source and target volumes is that the assumption that there exists a 1-to-1 homology between the two brains can no longer be validated. Even with this limitation, the images in Fig. 8 show that ANIMAL is able to recover the global brain shape and match individual brain structures. One can see in the difference images that there is almost perfect registration of the central portion of the brain (around the basal ganglia and ventricles) — precisely where the assumption of homology is valid. The region of the cortex where the difference images are non-null corresponds to regions where precise homology does not exist for all structures. Even though it is not clear what the correct registration is, nor even if there exists such a concept for non-homologous structures, the images of Fig. 9 shows that ANIMAL is able to use the labels from the target volume to segment the source.

It is important to note that the notion of 1-to-1 correspondence depends on the spatial scale of comparison. When the source and target datasets are blurred at 16 mm FWHM, all structures of the brain are topologically equivalent because only major structures (temporal lobe, ventricles, longitudinal fissure) are apparent. Smaller scale structures are simply not visible and do not take part in the matching process. When ANIMAL is applied to this scale of data, the overall shape of the head and brain are corrected. At the next scale, data is blurred with a Gaussian kernel with an 8 mm FWHM. ANIMAL is still able to register structures as long as there is a 1-to-1 mapping between them such as the ventricles, basal ganglia and major sulci. However, problems occur on the cortex, where topology is not consistent for secondary and tertiary gyri.[31] This is a complex problem beyond the scope of this paper and has been treated by other authors.[32,33]

There are a number of directions we plan to follow in the near future to improve the algorithm. Presently, ANIMAL uses only the gradient magnitude feature when computing the similarity between two neighborhood regions. To improve registration at the cortex, we have been looking at other features such as the Lvv described

by Florack *et al.*,[34] or explicitly extracted sulci.[35] Along the same lines, we are also looking at other new promising similarity metrics such as the mutual information criterion described by Collignon[36] and comparing them against the correlation function used in the current implementation. Finally, the variability map estimated in Sec. 4.3 provides information that can be used to constrain the estimation process on other data volumes. We are looking at probabilistic methods to incorporate this information into the ANIMAL algorithm.

6. CONCLUSIONS

We have presented a completely automatic method for non-linear registration of volumetric MRI data of the human brain. The design of the ANIMAL procedure meets the two goals set forth at the beginning of this paper, namely that of automatic segmentation and automatic estimation of anatomical variability, when using a pre-labeled target volume, resident in a standardized coordinate system. The multiresolution nonlinear registration strategy has been shown to be accurate and robust on both simulated and real MRI data, where it was able to recover important deformations, of up to 20 mm. The algorithm is useful for automatic labeling of individual brain structures in newly-acquired MRI volumes, by matching a previously-labeled MRI volume to the new unlabeled MRI volume. Thus, the redefinition of segmentation as a registration problem has been shown to be valid. The experimental results also demonstrate the practical utility of an automated 3-D image deformation algorithm for quantifying and removing, if so desired, anatomical variability among individual brains.

ACKNOWLEDGMENTS

The authors would like to express their appreciation for support from the Human Frontier Science Project Organization, the Canadian Medical Research Council (SP-30), the McDonnell-Pew Cognitive Neuroscience Center Program, the U.S. Human Brain Map Project (HBMP), NIMH and NIDA. This work forms part of a continuing project of the HBMP-funded International Consortium for Brain Mapping (ICBM) to develop a probabilistic atlas of human neuroanatomy. We also wish to acknowledge the manual anatomical labeling completed by Colin Holmes and Noor Kabani.

REFERENCES

1. P. T. Fox, M. A. Mintun, E. M. Reiman and M. E. Raichle, "Enhanced detection of focal brain responses using intersubject averaging and change-distribution analysis of subtracted PET images", *J. Cerebral Blood Flow and Metabolism* **8** (1988) 642–653.
2. P. T. Fox, J. S. Perlmutter and M. E. Raichle, "A stereotactic method of anatomical localization for positron emission tomography", *J. Comput. Assist. Tomogr.* **9**, 1 (1985) 141–153.
3. J. Talairach and P. Tournoux, *Co-planar Stereotactic Atlas of the Human Brain: 3-Dimensional Proportional System: An Approach to Cerebral Imaging*, Georg Thieme Verlag, Stuttgart, New York, 1988.

4. A. Evans, M. Kamber, D. Collins and M. D. Collins, "An MRI-based probabilistic atlas of neuroanatomy", *Magnetic Resonance Scanning and Epilepsy*, eds. S. Shorvon, D. Fish, F. Andermann, G. Bydder, and H. Stefan, vol. 264 of *NATO ASI Series A, Life Sciences*, Plenum Press, 1994, pp. 263–274.
5. A. C. Evans, D. L. Collins and B. Milner, "An MRI-based stereotactic atlas from 250 young normal subjects", *Soc. Neurosci. Abstr.* **18** (1992) 408.
6. C. Sorlié, D. L. Collins, K. J. Worsley and A. C. Evans, "An anatomical variability study based on landmarks", *Human Brain Mapping*, Technical Report, 1994.
7. F. Bookstein, "Thin-plate splines and the atlas problem for biomedical images", *Information Processing in Medical Imaging*, eds. A. Colchester and D. Hawkes, vol. 511 of *Lecture Notes in Computer Science*, IPMI, Springer-Verlag, Wye, UK, July 1991, pp. 326–342.
8. A. C. Evans, W. Dai, D. L. Collins, P. Neelin and T. Marrett, "Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis", *Proc. Int. Soc. Optical Engineering: Medical Imaging V*, vol. 1445, SPIE, San Jose, California, 27 Feb.–1 Mar. 1991.
9. Y. Ge, J. Fitzpatrick, R. Kessler and R. Margolin, "Intersubject brain image registration using both cortical and subcortical landmarks", *Proc. SPIE Medical Imaging* **2434** (1995) 81–95.
10. S. Luo and A. Evans, "Matching sulci in 3-D space using force-based deformation", *IEEE Trans. Med. Imag.*, Technical Report, 1994.
11. N. A. G. Subsol and J.-P. Thirion, "Application of an automatically built 3-D morphometric brain atlas: Study of cerebral ventricle shape", *Proc. Conf. Visualization in Biomedical Computing*, Lecture Notes in Computer Science, Springer-Verlag, Sept. 1996, pp. 373–382.
12. J. Downs III, L. J. L. and F. P. T., "3-D surface based spatial normalization using a convex hull," *Functional Neuroimaging: Technical Foundations*, eds. R. Thatcher, M. Hallett, T. Zeffiro, E. John and M. Huerta, 1994, pp. 63–80.
13. J. Downs III, L. J.L., and F. P.T., "Intersubject variability of cerebral structures after spatial normalization using a cortical convex hull", *Soc. Neurosci. Abstr.* **1**, 354 (1994) 152.61.
14. P. Thompson and A. Toga, "A surface-based technique for warping 3-dimensional images of the brain," *IEEE Trans. Med. Imag.* **15**, 4 (1996) 1–16.
15. S. Sandor and R. Leahy, "Towards automated labeling of the cerebral cortex using a deformable atlas", *Information Processing Medical Imaging*, eds. Y. Bizais, C. Barillot and R. DiPaola, Brest, France, IPMI, Kluwer, Aug. 1995, pp. 127–138.
16. C. Broit, *Optimal Registration of Deformed Images*, PhD thesis, University of Pennsylvania, Philadelphia, 1981.
17. R. Bajcsy and C. Broit, "Matching of deformed images", *Proc. 6th Int. Conf. Pattern Recognition*, IEEE, Munich, Germany, Oct. 19–22 1982, pp. 351–353.
18. R. Bajcsy and S. Kovacic, "Multiresolution elastic matching", *Comput. Vision Graphics Imag. Process.* **46** (1989) 1–21.
19. R. Dann, J. Hoford, S. Kovacic, M. Reivich and R. Bajcsy, "Three-dimensional computerized brain atlas for elastic matching: Creation and initial evaluation", *Medical Imaging II*, SPIE, Newport Beach, California, Feb. 1988, pp. 600–608.
20. J. Gee, "Probabilistic matching of deformed images", Tech. Rep. Technical report MS-CIS-96, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, 1996.
21. J. Gee, L. LeBriquer and C. Barillot, "Probabilistic matching of brain images", *Information Processing Medical Imaging*, eds. Y. Bizais and C. Barillot, IPMI, Kluwer, Berder, France, July 1995.

22. M. Miller, Y. A. G. E. Christensen and U. Grenander, "Mathematical textbook of deformable neuroanatomies", *Proc. National Academy of Sciences* **90**, 24 (1990) 11944–11948.
23. G. Christensen, R. Rabbitt and M. Miller, "3D brain mapping using a deformable neuroanatomy", *Phys. Med. Biol.* **39** (1994) 609–618.
24. G. Christensen, R. Rabbitt and M. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. Image Process.* **5**, 10 (1996) 1435–1447.
25. S. Joshi, M. Miller, G. Christensen, A. Banerjee, T. Coogan and U. Grenander, "Hierarchical brain mapping via a generalized dirichlet solution for mapping brain manifolds", *Proc. SPIE's 1995 Int. Symp. Optical Science, Engineering, and Instrumentation*, vol. 2573, *Vision Geometry IV*, SPIE, Aug. 1995, pp. 278–289.
26. J. Ashburner and K. Friston., "Fully three-dimensional nonlinear spatial normalisation: A new approach", *2nd Int. Conf. Functional Mapping of the Human Brain*, eds. J. Belliveau, D. Kennedy, and B. Rosen, page accepted, Organization for Human Brain Mapping, Boston, June 1996.
27. K. Friston, , J. Ashburner, C. Frith, J.-B. Poline, J. Heather and R. Frackowiak, "Spatial registration and normalization of images", *Human Brain Mapping* **1**, 2 (1995) 1–25.
28. J. Koenderink and A. van Doorn, "Representation of local geometry in the visual system", *Biol. Cybern.* **55** (1987) 367–375.
29. B. tar Haar Romeny, L. M. Florack, J. J. Koenderink and M. A. Viergever, "Scale space: Its natural operators and differential invariants", *Information Processing Medical Imaging*, eds. A. C. F. Colchester and D. J. Hawkes IPMI, Wye, UK, July 1991, p. 239.
30. D. L. Collins, P. Neelin, T. M. Peters and A. C. Evans, "Automatic 3D inter-subject registration of MR volumetric data in standardized talairach space", *J. Comput. Assist. Tomogr.* **18** (1994) 192–205.
31. M. Ono, S. Kubik and C. Abernathey, *Atlas of Cerebral Sulci*, Stuttgart: Georg Thieme Verlag, 1990.
32. D. MacDonald, D. Avis and A. C. Evans, "Multiple surface identification and matching in magnetic resonance images", *Proc. Conf. Visualization in Biomedical Computing*, SPIE, 1994.
33. J.-F. Mangin, V. Frouin, I. Bloch, J. Regis and J. Lopez-Krahe, "From 3d magnetic resonance images to structural representations of the cortex topology using topology preserving deformations", *J. Mathematical Imaging and Vision*, accepted, 1995.
34. L. M. J. Florack, B. M. ter Haar Romeny, J. J. Koenderink and M. A. Viergever, "Scale and the differential structure of images", *Image Vision Comput.* **10** (1992) 376–388.
35. D. Collins, G. LeGoualher, R. Venugopal, Z. Caramanos, A. Evans and C. Barillot, "Cortical constraints for nonlinear cortical registration" *Visualization in Biomedical Computing*, ed. K. Höene, Hamburg, Sept. 1996, pp. 307–316.
36. A. Collingnon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens and G. Marchal, "Automated multi-modality image registration based on information theory", *Proc. Information Processing in Medical Imaging*, Brest, France, June 1995, pp. 263–274.

Louis Collins received the Bachelor's degree in computer science from Concordia University (Montreal) in 1987, a Master's degree in electrical engineering and a PhD in biomedical engineering from McGill University (Montreal). After completion of an HFSPO-funded post-doctoral fellowship at the Laboratoire S.I.M. at the University of Rennes in France, he now holds a position as a Research Associate with the McConnell Brain Imaging Centre of the Montreal Neurological Institute.

His research interests include inter-subject medical image registratation, and model-based segmentation to automatically identify structures within the human brain.

Dr. Evans completed his Ph.D. and post-doctoral fellowship studying structure-function interaction of proteins at the Department of Biophysics at Leeds University. He moved to Atomic Energy of Canada Ltd. as a PET physicist interested in imaging of brain function. He joined the department of Neurology/Neurosurgery at McGill University in 1984, where he became a full professor in 1995 with appointments in Neurology/Neurosurgery and Biomedical Engineering. He co-founded the NeuroImaging Laboratory at the Montreal Neurological Institute (MNI) in 1988. He has been Coordinator of the MNI's McConnell Brain Imaging Center since 1992.

His research interests include cognitive neuroimaging, neuroanatomical variability, and image processing methodologies for PET and MRI.