

HAMMER: Hierarchical Attribute Matching Mechanism for Elastic Registration

Dinggang Shen* and Christos Davatzikos, *Member, IEEE*

Abstract—A new approach is presented for elastic registration of medical images, and is applied to magnetic resonance images of the brain. Experimental results demonstrate very high accuracy in superposition of images from different subjects. There are two major novelties in the proposed algorithm. First, it uses an attribute vector, i.e., a set of geometric moment invariants (GMIs) that are defined on each voxel in an image and are calculated from the tissue maps, to reflect the underlying anatomy at different scales. The attribute vector, if rich enough, can distinguish between different parts of an image, which helps establish anatomical correspondences in the deformation procedure; it also helps reduce local minima, by reducing ambiguity in potential matches. This is a fundamental deviation of our method, referred to as the hierarchical attribute matching mechanism for elastic registration (HAMMER), from other volumetric deformation methods, which are typically based on maximizing image similarity. Second, in order to avoid being trapped by local minima, i.e., suboptimal poor matches, HAMMER uses a successive approximation of the energy function being optimized by lower dimensional smooth energy functions, which are constructed to have significantly fewer local minima. This is achieved by hierarchically selecting the driving features that have distinct attribute vectors, thus, drastically reducing ambiguity in finding correspondence. A number of experiments demonstrate that the proposed algorithm results in accurate superposition of image data from individuals with significant anatomical differences.

Index Terms—Attribute vectors, average brain, deformable registration, geometric moment invariants, hierarchical deformation mechanism, multigrid formulation, statistical atlases.

I. INTRODUCTION

DEFORMABLE registration of brain images has been an active topic of research for over a decade. Its clinical applications are numerous. In particular, deformable registration is used for spatial normalization of functional images, group analysis, and statistical parametric mapping [1]. It is also used in computational anatomy as a means for measuring structures,

by adapting an anatomical template to individual anatomies [2]–[11], [46]. Finally, it is used as a means for image data mining in lesion-deficit studies [12], as well as in stereotaxic neurosurgery for mapping anatomical atlases to patient images [13]–[16]. Therefore, many image analysis methodologies have been developed to tackle this issue, which fall in two general categories. The first family of methods involves feature-based matching, i.e., transformations that are calculated based on a number of anatomical correspondences established manually, semiautomatically, or fully automatically on a number of distinct anatomical features. Such features are distinct landmark points [17], [18], or a combination of curves and surfaces, such as sulci and gyri [6], [19]–[24]. The second family of methods is based on volumetric transformations, which seek to maximize the similarity between an image and a template, and generally assume that the image and the template have been acquired with the same imaging protocol [1], [25]–[30], [47], [48], [51].

Each of these approaches has its advantages and disadvantages. Feature-based methods utilize anatomical knowledge in determining point correspondences, and can be faster, since they do not evaluate a matching criterion on every single voxel in an image, but they rely on a relatively small number of feature points. Similarity-based methods can be fully automated, and are more general, since they do not require the construction of a specific anatomical model each time they are applied to a new problem. However, they do not directly solve the problem of anatomical correspondences, as image similarity does not necessarily imply good registration of the underlying anatomy. Moreover, since gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF) have practically uniform image intensities throughout the whole brain, optimality criteria based on image similarity can suffer significantly from local minima, which are caused by the ambiguity in defining correspondence.

In medical imaging, it is important to build deformable anatomical models that take into account the underlying anatomy [51], and not simply the similarity of image intensity. Toward that end, we previously introduced an adaptive focus deformable model (AFDM) [31], which utilized the concept of an *attribute vector*, i.e., a vector of geometric attributes that was attached to each point on a surface model of an anatomical structure, and which reflected the geometric properties of the underlying structure from a local scale (e.g., curvature), to a global scale that reflected spatial relationships with more distant surface points. If the attribute vector is rich enough, it can differentiate between different parts of the anatomy that would otherwise look similar, if only image characteristics were taken into account (e.g., two points that lie on the hippocampal boundary, but belong to different parts of the structure) [32].

Manuscript received July 26, 2001; revised May 9, 2002. This work was supported in part by the National Institute of Health (NIH) under Grant NIH-R01AG14971 and under Contract AG-93-07. Images were acquired as part of the neuroimaging study of the Baltimore Longitudinal Study of Aging [49]. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was M. Sonka. *Asterisk indicates corresponding author.*

*D. Shen was with the Center for Biomedical Image Computing, Department of Radiology, The Johns Hopkins University School of Medicine, 601 N. Caroline Street, Baltimore, MD 21287 USA. He is now with the Department of Radiology, University of Pennsylvania School of Medicine, Philadelphia, PA 19104 USA (e-mail: dgshen@rad.upenn.edu).

C. Davatzikos was with the Center for Biomedical Image Computing, Department of Radiology, The Johns Hopkins University School of Medicine, Baltimore, MD 21287 USA. He is now with the Department of Radiology, University of Pennsylvania School of Medicine, Philadelphia, PA 19104 USA.

Digital Object Identifier 10.1109/TMI.2002.803111

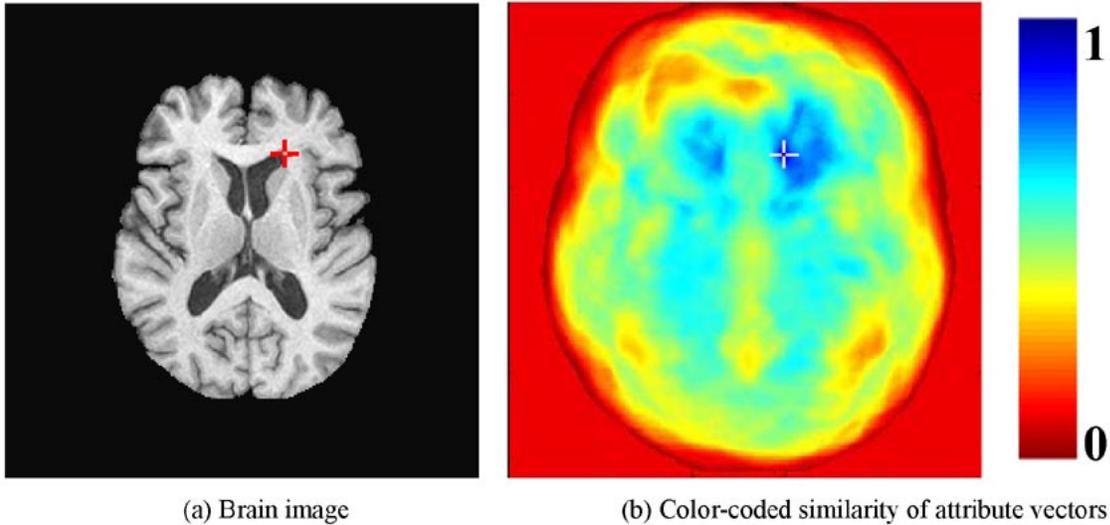


Fig. 1. Demonstration of how GMIs distinguish one brain region from others. (a) Brain image, where the GMIs of the WM voxel indicated by a red cross are compared with the GMIs of all other voxels. The size of this 2-D image is 256×256 , and the GMIs are calculated in a 81×81 neighborhood. The resulting similarity of the GMIs is shown color-coded in (b), with blue reflecting high similarity. The white cross in (b) corresponds to the red cross in (a). This figure demonstrates that the attribute vectors can distinguish one voxel from other voxels that could have similar intensity or edge type.

In this paper, we build upon the ideas of the approach in [31], by defining attribute vectors in volumetric images, and using them for three-dimensional (3-D) warping. Our focus is on minimizing the effect of local minima of the energy function being optimized, and on determining anatomically meaningful image correspondence. The presence of local minima is due primarily to three reasons: 1) the very high dimensionality of the space formed by the coordinates of all voxels in a volumetric image; 2) ambiguities inherent in the matching function being maximized and related to the practically uniform image intensities of GM, WM, and CSF throughout the brain; and 3) the complex nature of brain anatomy. We address the first two issues above as follows.

1) We use a sequence of successive approximations of the energy function being optimized, by lower dimensional smooth energy functions. This is accomplished by a hierarchical procedure, which initially uses only a small set of features to drive the volumetric deformation and gradually increases the number of features and, therefore, increases the dimensionality of the energy function. The driving features are selected hierarchically according to how uniquely they can be identified among other features in the image, which reduces ambiguity and therefore local minima. This introduces the second novelty of the algorithm, described next.

2) In order to characterize the brain structure in the vicinity of each voxel, we use the concept of an attribute vector, which is defined on each voxel in an image, and which reflects the underlying anatomy at different scales. In addition to image intensity and edge information, this attribute vector includes geometric moment invariants (GMIs) as a means for representing the geometric structure of the underlying anatomy. GMIs have been used very successfully in different computer vision applications [33]–[37], [45] in order to represent the local structure of images. The idea behind the attribute vector is that, if it is rich enough, it can distinguish between different parts of

an image, which might otherwise be indistinguishable. For example, all GM voxels have similar intensities. Therefore, when trying to find a spatial transformation that maximizes image similarity, the precentral gyrus might well be matched to its neighboring postcentral gyrus, which has the same image intensity. However, when examined at different scales, the precentral and postcentral gyri have different characteristics, hence, anatomists have assigned them different names. GMIs at a particular scale are calculated by placing a spherical neighborhood around each voxel, and then calculating a number of parameters that are invariant to rotation [38]. The radius of the neighborhood defines the scale of the GMIs. We evaluate the GMIs at different scales (i.e., for different neighborhood sizes), and concatenate the resulting GMIs into a long attribute vector. The idea of GMIs is demonstrated schematically in Fig. 1 by a two-dimensional (2-D) image, where different local structures have different GMIs.

We present experiments with images from elderly subjects, which display several difficulties, including reduced tissue contrast, significant atrophy, and motion artifacts. We demonstrate patient results with high qualitative accuracy (see Figs. 12 and 13), and results from simulations showing the average registration error of 0.966 mm.

II. METHODS

A. General Formulation

Let us assume that $T(x)$ is the intensity of the voxel x in the template brain image, and $S(y)$ is the intensity of the voxel y in the individual's brain image. Here, x and y are the 3-D coordinates of the voxels, respectively, in the volumes V_T and V_S . The displacement field $u(x)$ defines the mapping from the coordinate system of the template T to the subject S , while the transformation $h(x) = x + u(x)$ defines the mapping that transforms the template T into the shape of the subject S . The inverse

transformation of $h(x)$ is $h^{-1}(x)$. By transforming the template to the subject, for each voxel in the template we know its corresponding voxel in the subject. In this way, we can warp the subject to the space of the template. Therefore, our goal here is to warp the subject to the template by deforming the template to the subject.

The hierarchical attribute matching mechanism for elastic registration (HAMMER) uses a sequence of lower dimensional energy functions to ultimately approximate the following multivariate energy function:

$$E = \sum_{x \in V_T} \omega_T(x) \left(\frac{\sum_{z \in n(x)} \varepsilon(z) (1 - m(a_T(z), a_S(h(z))))}{\sum_{z \in n(x)} \varepsilon(z)} \right) + \sum_{y \in V_S} \omega_S(y) \left(\frac{\sum_{z \in n(y)} \varepsilon(z) (1 - m(a_T(h^{-1}(z)), a_S(z)))}{\sum_{z \in n(y)} \varepsilon(z)} \right) + \beta \sum_x \|\nabla^2 u(x)\|. \quad (1)$$

The first energy term is the summation of similarity matching in the template space, for calculating the transformation $h(\cdot)$ from the template to the subject. For each template voxel x , we compute the similarity of the attribute vectors in its spherical neighborhood $n(x)$, rather than on individual voxels. For example, for a neighboring template voxel z , we compare the attribute vector of this template voxel z with the attribute vector of its corresponding deformed voxel $h(z)$ in the subject. $\mathbf{a}_T(z)$ is the attribute vector of the template voxel z , while $\mathbf{a}_S(h(z))$ is the attribute vector of the subject voxel at the position $h(z)$. The function $m(\cdot, \cdot)$ measures the similarity of two attribute vectors, with the range from 0 to 1. Therefore, $1 - m(\cdot, \cdot)$ measures the difference of two attribute vectors. $\varepsilon(\cdot)$ is a weight, which assigns larger values to the boundary voxels, since those voxels are important anatomical features and are typically easy to find in the images of sufficient contrast. The term $\sum_{z \in n(x)} \varepsilon(z)$ is used for normalization. The weighting parameter $\omega_T(x)$ gives the relative weight for each voxel x in the template image. We assign larger weights to the voxels with the distinct attribute vectors, which usually can be identified relatively more reliably than others. These voxels will initially drive the deformation of the template to the subject. By hierarchically focusing on the different energy terms in this energy function (major terms first), our model is able to hierarchically deform the template to the subject.

The second energy term is similar to the first, but it is defined in the subject domain, and used to constrain the inverse transformation $h^{-1}(\cdot)$, from the subject to the template. It is a consistency of the transformation term, analogous to the formulation that was originally proposed in [42] and [43]. Since the weights $\omega_S(y)$ and $\omega_T(x)$ are determined by the image details, which are different in the template and the subject, the emphases of these two energy terms are different. The first energy term focuses on deforming “template to subject,” while the second energy term focuses on pulling “subject to template.” This formulation makes HAMMER robust to initialization and suboptimal solutions (see examples in Fig. 7). Naturally, the introduction of this

term increases the computational requirements of our model. In order to overcome this, we selectively focus on only a small set of the subject voxels, which we call the *driving voxels*, i.e., the voxels of the distinct attribute vectors that are described later in this section. Although this does not ensure consistency of the full 3-D transformation [42], [43], it ensures consistency of the driving forces. Since the driving forces are what effectively determines correspondence between two images, and are simply interpolated elsewhere in the volume, this approach has turned out to be adequate in our experiments.

The third term is a smoothness constraint on the displacement fields. Here, ∇^2 is Laplacian operator. This term is similar to the regularization terms that were extensively used in other literature [17], [52]. The parameter β controls the smoothness of the deformation fields. But the large β might lead to oversmoothness of the deformation fields. For our application, we used 0.5 for β .

B. Attribute Vectors

As described earlier, the attribute vectors are a critical part of our formulation. The attribute vector is defined for each voxel x in a volumetric image, and it reflects the underlying structure at different scales. Each attribute vector includes edge types, image intensity, and GMIs. If rich enough, the attribute vector can distinguish between different parts of the image. The attribute vector $\mathbf{a}(x)$ at the voxel x is made up of three vectors $\mathbf{a}_1(x)$, $\mathbf{a}_2(x)$, and $\mathbf{a}_3(x)$. That is

$$\mathbf{a}(x) = [\mathbf{a}_1(x) \quad \mathbf{a}_2(x) \quad \mathbf{a}_3(x)].$$

$\mathbf{a}_1(x)$ is a 1×1 vector. It represents the edge type of the voxel x in the image. For the hard tissue segmentation case [39] that we are using in this paper, it takes one of seven discrete values, corresponding to seven edge types such as nonedge and six combinations of edges between GM, WM, and CSF. For the fuzzy segmentation case [50], each voxel x has the tissue classification result such as a vector $[\mathbf{C}_{WM}, \mathbf{C}_{GM}, \mathbf{C}_{CSF}]^T$, where \mathbf{C}_{WM} , \mathbf{C}_{GM} , and \mathbf{C}_{CSF} take the real values between 0 and 1.0, and $\mathbf{C}_{WM} + \mathbf{C}_{GM} + \mathbf{C}_{CSF} = 1.0$. Then, $\mathbf{a}_1(x)$ can be designed as a 3×1 vector, such as the direct difference (3×1 vector) of the tissue classification results between the voxel x and its neighboring voxel z , where the voxel z makes the magnitude of this difference maximal in all neighboring voxels.

$\mathbf{a}_2(x)$ is a 1×1 vector (but in general could have more components), and it represents the intensity of the voxel x , with the range of the gray levels. For the convenience, $\mathbf{a}_2(x)$ is normalized between zero and one.

The vector $\mathbf{a}_3(x)$ comprises the GMIs of each tissue at different scales. In the experiments herein, we have used hard classification to label brain tissue into the three types GM, WM, and CSF. For each scale and each tissue type, there are 13 rotation invariants that are calculated from the zero-order, second-order, and third-order 3-D regular moments [38]. Therefore, $\mathbf{a}_3(x)$ is a $3 \times 13 \times L$ vector, where L is the number of scales used. In the following, we provide the mathematical definition of GMIs. For convenience, let us assume that the origin of the coordinate system has been shifted to the voxel x under consideration, and that under this coordinate system, the membership function (either GM, WM, or CSF), is $f_{\text{tissue}}(x_1, x_2, x_3)$,

where x_1 , x_2 , and x_3 are the coordinates of a voxel. Then, the 3-D regular moments of the order $(p + q + r)$ of the function $f_{\text{tissue}}(x_1, x_2, x_3)$ are defined by

$$M_{p, q, r} = \iiint_{(x_1)^2 + (x_2)^2 + (x_3)^2 < R^2} x_1^p x_2^q x_3^r f_{\text{tissue}}(x_1, x_2, x_3) dx_1 dx_2 dx_3$$

where R is the radius of the spherical neighborhood around the origin (i.e., the voxel x). Four rotation invariants that are formulated from the zero-order and the second-order moments are listed next. (An additional nine rotation invariants, which are formulated from the third-order moments and from both of the second-order and the third-order moments, are omitted here, due to their very complicated expressions. They can be found in [38].)

$$\begin{aligned} I_1 &= M_{0,0,0} \\ I_2 &= M_{2,0,0} + M_{0,2,0} + M_{0,0,2} \\ I_3 &= M_{2,0,0}M_{0,2,0} + M_{2,0,0}M_{0,0,2} + M_{0,2,0}M_{0,0,2} \\ &\quad - M_{1,0,1}^2 - M_{1,1,0}^2 - M_{0,1,1}^2 \\ I_4 &= M_{2,0,0}M_{0,2,0}M_{0,0,2} - M_{0,0,2}M_{1,1,0}^2 \\ &\quad + 2M_{1,1,0}M_{1,0,1}M_{0,1,1} - M_{0,2,0}M_{1,0,1}^2 \\ &\quad - M_{2,0,0}M_{0,1,1}^2. \end{aligned}$$

Notice here that the first GMI corresponds to the volume of a certain tissue.

GMI is a convenient way of characterizing the geometric properties of objects in 2-D or 3-D images, since objects of different shapes tend to have different GMIs. Fig. 1 demonstrates the concept of GMIs that can be used to discriminate local structures. For the convenience, we used a 2-D image as an example. For the WM voxel that was indicated by a red cross in Fig. 1(a), there are many other WM voxels in its neighborhood. Using the image intensity-matching mechanism, we cannot distinguish this voxel from the others. But using the attribute vector, we can do this. Fig. 1(b) shows a color-coded map of the similarity between the attribute vector of this WM voxel and the attribute vectors that are calculated in other positions. The similarity degree ranges from zero to one, and blue is used for high similarity. From this map, we can observe that this voxel is only similar to its neighboring voxels, and to some extent to its symmetric voxels. It should be indicated here that besides using GMIs to extract the geometric properties of objects, other techniques such as Gabor filters, wavelets, and low-frequency representations [41] can also be applied, *under the condition that the extracted attributes are rotation invariant*. The property of rotation invariance is very important; a sulcus or gyrus should be detected as such, regardless of its orientation in the image.

Having defined the attribute vector, we now define the similarity criterion used for deforming the template to the subject. We require that boundary voxels in the template deform to the voxels with the same boundary type in the subject, since boundaries are very important features in describing brain structure. This also reduces computational requirements. By using this similarity criterion, we can make the warped subject's boundaries match well with the corresponding template boundaries

(cf. the right corner of Fig. 9). On the other hand, this similarity criterion is robust to tissue segmentation errors [39], because our deformation strategy integrates the similarity of attribute vectors within a subvolume, instead of evaluating it on an individual voxel. In fact, most of the images in this paper are from elderly subjects, who tend to display signal abnormalities that result in misclassification of some WM as GM. These errors, however, have not affected the performance of HAMMER.

We define the similarity of the two voxels x and y as

$$m(\mathbf{a}(x), \mathbf{a}(y)) = \begin{cases} 0, & \text{if } \mathbf{a}_1(x) \neq \mathbf{a}_1(y) \\ c([\mathbf{a}_2(x) \ \mathbf{a}_3(x)], [\mathbf{a}_2(y) \ \mathbf{a}_3(y)]), & \text{otherwise} \end{cases}$$

where $c([\mathbf{a}_2(x) \ \mathbf{a}_3(x)], [\mathbf{a}_2(y) \ \mathbf{a}_3(y)])$ is the similarity of the second and third parts of the attribute vectors. When a set of training samples is available, the similarity $c([\mathbf{a}_2(x) \ \mathbf{a}_3(x)], [\mathbf{a}_2(y) \ \mathbf{a}_3(y)])$ can be learned from the training samples using learning techniques such as support vector machines [40]. Without a set of training samples, we simply normalize each element in \mathbf{a}_2 and \mathbf{a}_3 to a range from zero to one, and then define $c([\mathbf{a}_2(x) \ \mathbf{a}_3(x)], [\mathbf{a}_2(y) \ \mathbf{a}_3(y)])$ as

$$c([\mathbf{a}_2(x) \ \mathbf{a}_3(x)], [\mathbf{a}_2(y) \ \mathbf{a}_3(y)]) = (1 - |\mathbf{a}_2(x) - \mathbf{a}_2(y)|) \cdot \prod_{i=1}^K (1 - |\mathbf{a}_3^i(x) - \mathbf{a}_3^i(y)|)$$

where $\mathbf{a}_3^i(x)$ is the i th element of $\mathbf{a}_3(x)$ that has a total of K elements. This is the definition used throughout the experiments of this paper.

C. Image Deformation Mechanism: Hierarchical Approximation of the Energy Function

Overview: In order to optimize the criterion in (1), we developed a deformation mechanism that is robust to sub-optimal solutions. This deformation mechanism also results in displacement fields that are well behaved, in that they are smooth and have smooth first derivatives. This deformation mechanism is based on the approximation of the energy function $E = E(x_1^T, x_2^T, \dots, x_{P_T}^T; y_1^S, y_2^S, \dots, y_{P_S}^S)$ in (1), where $x_1^T, x_2^T, \dots, x_{P_T}^T$ are the total P_T voxels in the template volumetric image and $y_1^S, y_2^S, \dots, y_{P_S}^S$ are the total P_S voxels in the target volumetric image, by an energy function $\varepsilon = \varepsilon(x_1, x_2, \dots, x_{N_T}; y_1, y_2, \dots, y_{N_S})$, $N_T < P_T$ and $N_S < P_S$, where x_1, x_2, \dots, x_{N_T} are N_T driving voxels in the template and y_1, y_2, \dots, y_{N_S} are N_S driving voxels in the subject. Initially, $N_T \ll P_T$ and $N_S \ll P_S$, thus implying that only few distinctly identifiable voxels are used to drive the deformation. The displacements $\Delta(x_i^T)$ of all other voxels in the template, which are interpolated from the displacements of the driving voxels, are given by

$$\Delta(x_i^T) = \sum_{j=1}^{N_T} \Delta(x_j) g_{ij} \quad (2)$$

where $\Delta(x_j)$ is a displacement on the driving voxel x_j and g_{ij} is a Gaussian kernel which prescribes that the influence of the driving voxel x_j on the displacement of other voxels fades away with distance. In this way, neighboring voxels deform

smoothly and together with the driving voxel in a Gaussian way. This formulation is important, particularly for the initial stages, where the driving voxels are sparse in the space. Equation (2) readily leads to the approximation of the full energy function $E = E(x_1^T, x_2^T, \dots, x_{P_T}^T; y_1^S, y_2^S, \dots, y_{P_S}^S)$ by a function $\varepsilon = \varepsilon(x_1, x_2, \dots, x_{N_T}; y_1, y_2, \dots, y_{N_S})$. In order to implement (2) in a practically feasible fashion, in all experiments presented herein we subdivide the volumetric image into discrete *subvolumes*, each of which is effectively a volume of influence around a driving voxel, with the appropriate continuity boundary constraints. Therefore, the displacement of each template voxel is influenced only by a single driving voxel, which makes the sum in (2) a single term. In other words, we propagate the displacement of a driving voxel to a subvolume around it, according to a Gaussian kernel, depending on the distance from the driving voxel. This deformation mechanism is described in more detail later in this section.

The energy function $\varepsilon(x_1, x_2, \dots, x_{N_T}; y_1, y_2, \dots, y_{N_S})$ has significantly fewer local minima than $E(x_1^T, x_2^T, \dots, x_{P_T}^T; y_1^S, y_2^S, \dots, y_{P_S}^S)$ for two reasons. First, it is based on a significantly lower number of variables. Second, these variables are selected to be relatively unambiguously identified features, as described below.

Driving Voxels: Some parts of the anatomy can be identified more reliably than others. This is due to several reasons. First, some parts have very distinct geometric characteristics, which identify them indisputably. Good examples are the roots of sulci and the crowns of the gyri, which can be identified much more reliably than intermediate cortical points. Second, the attribute vectors, which in our framework provide the means by which anatomical characteristics are represented, might be better discriminants of some features over others. Finally, certain parts of the anatomy are more complex than others. For example, in regions where many structures are close to each other, edges extracted from the images are dense, and might simultaneously influence the deformation of deformable model, thus rendering its deformation prone to errors. We address exactly these issues in our approach to determining driving voxels. Specifically, we use the following rules.

1) Roots of sulci and crowns of gyri are typically identified very robustly, based on their distinctive attribute vectors. This is demonstrated in Fig. 3(a1). Accordingly, the deformation of the model is influenced primarily by those regions initially, and gradually shifts focus to other cortical regions, as those get close to their respective targets. This hierarchical scheme is implemented by focusing on a certain range of attribute vector values at a particular stage of the deformation. This range is initially relatively narrow, thus including only sulcal roots and gyral crowns, and increases gradually as the model converges.

2) Regions that display relatively high similarity of the respective attribute vectors have a relatively stronger influence on the deformation mechanism. This characteristic makes the deformation mechanism robust to spurious matches, at least in the beginning of the deformation, when the chances for the algorithm to be trapped in local minima are relatively high.

3) Voxels located on strong and isolated edges are usually easy to determine. These types of voxels can be detected by evaluating the difference between their edge strength and the mean

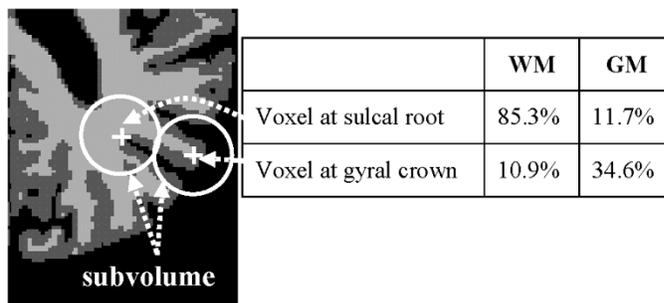


Fig. 2. Schematic explanation of why the first GMI of the WM is distinct for the sulcal roots and the gyral crowns. The table shows the percents of the WM and GM, respectively, in the spherical subvolumes of the two voxels at the sulcal root and the gyral crown, as indicated by two white crosses.

edge strength in their vicinities. In the future, we are planning to determine such voxels from a training set, in a way analogous to the approach presented in [31].

There are two reasons that cause the roots of sulci and the crowns of gyri to have distinctive attribute vectors. First, they are in the regions with high boundary curvatures and, therefore, their attribute vectors are very different from the intermediate cortical voxels. Here is a simple example. The neighborhoods around voxels in the roots of sulci usually include relatively large WM volume, which is represented by the first GMI of the WM, while the neighborhoods of voxels in the crowns of gyri include relatively small WM volume (see Fig. 2). Second, the number of the voxels that belong to roots of sulci and crowns of gyri is much smaller than the number of other brain voxels, which makes them fairly distinct, in that few other points that are nearby are likely to be good candidate matches.

The driving voxels are hierarchically selected in our algorithm in the following ways. We use a fuzzy clustering method to find three cluster centers, corresponding to the center of gyri, the center of sulci, and the center of the other boundary voxels, respectively. Initially, the voxels whose attribute vectors are close to the center of the gyri and the center of the sulci are selected as the driving voxels. With the increase of the iterations, the voxels whose attribute vectors are far from the center of the gyri and the center of the sulci are gradually added as the driving voxels. In the final stages, all voxels in the image domain are selected as the driving voxels. Since in this paper we apply HAMMER on the images from the elderly subjects that display significant ventricular atrophy, the ventricular boundary voxels are typically used as the initial driving voxels.

Fig. 3 provides an example of the driving voxels, which are hierarchically selected at the successive deformation stages that deform the template in Fig. 3(a1)–(a4) to the subject in Fig. 3(b). The subject image in Fig. 3(b) is the same as the cross-section image in Fig. 9(Li), and it is an affine-transformed version of the image in Fig. 9(I), with the affine transform automatically determined by HAMMER. The yellow points in Fig. 3(a1) are the voxels that our model automatically determined during the initial deformation stages. Most of them are the voxels of sulcal roots and gyral crowns. (Only voxels on the GM/WM interface are shown; hence, we are using the terms sulcal root and gyral crown loosely here.) After using these voxels to drag the major structures of the template to their corresponding parts in the sub-

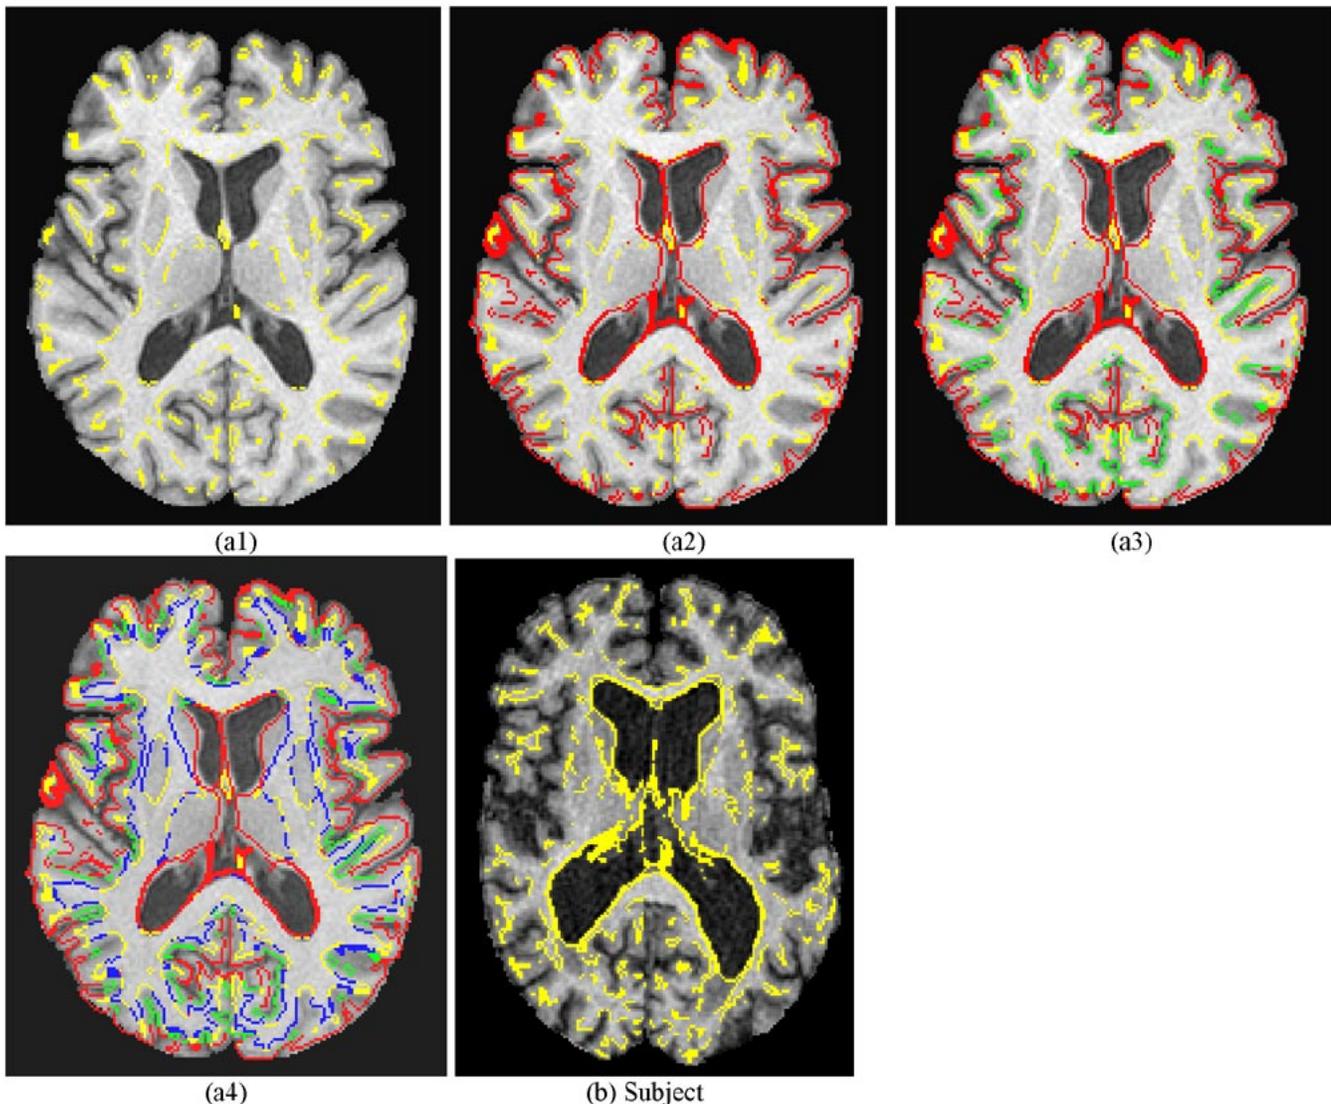


Fig. 3. Demonstration of the hierarchical deformation mechanism employed by HAMMER. (a1)–(a4) shows the driving voxels in the template at different deformation stages. Most of the template driving voxels selected in the initial deformation stage, which are shown as yellow in (a1), are the voxels in the sulcal roots and the gyral crowns. (Only voxels on the GM/WM interface are shown.) In the subsequent stages, more driving voxels are added, as the red voxels in (a2), the green voxels in (a3), and the blue voxels in (a4). The yellow voxels in (b) denote the driving voxels in the subject image, which are automatically determined and used to pull the template to the subject. Only voxels on the GM/WM interface are shown, for clarity.

ject, other boundary voxels are added to the set of the driving voxels and immediately join the procedure of deforming the template. For example, in the subsequent stages, the red voxels are added in Fig. 3(a2), then the green voxels are added in Fig. 3(a3), and finally, the blue voxels are added in Fig. 3(a4). The yellow voxels in Fig. 3(b) are the driving voxels in the subject image that have distinctive attribute vectors. Only voxels on the WM/GM boundary are shown, for clarity. It is notable that the blue points include many intermediate cortical points and many of the sub-cortical GM/WM boundary points, which are often difficult to identify accurately due to decreased tissue contrast.

Deformation of Subvolumes: Many deformable registration methods rely on iterative methods that are employed for the solution of some large and sparse linear or nonlinear system of equations, and which effectively move one voxel at a time, trying to improve some similarity criterion. This tendency has been largely inherited from standard numerical analysis

methods for solving such systems. However, the huge dimensionality of systems arising in 3-D warping methods of medical images, in conjunction with the highly nonconvex nature of the underlying energy functions resulting from the complex nature of brain anatomy, renders such iterative methods susceptible to local minima. Moreover, the resulting displacement fields are often “bumpy,” as they pull and push individual voxels to nearby voxels of similar intensity characteristics. Therefore, it is reasonable to expect that deforming larger regions of the image would result in more robust deformation schemes. This strategy has been adopted in B-spline-based models [52], which warp all voxels within the small neighborhood of a control point, if this control point is translated.

Similarly, in our approach, for each template voxel that is considered at a particular stage of the deformation (e.g., initially a voxel belonging to a sulcal root or a gyral crown), a subvolume, i.e., a small volume around that voxel, is displaced and deformed

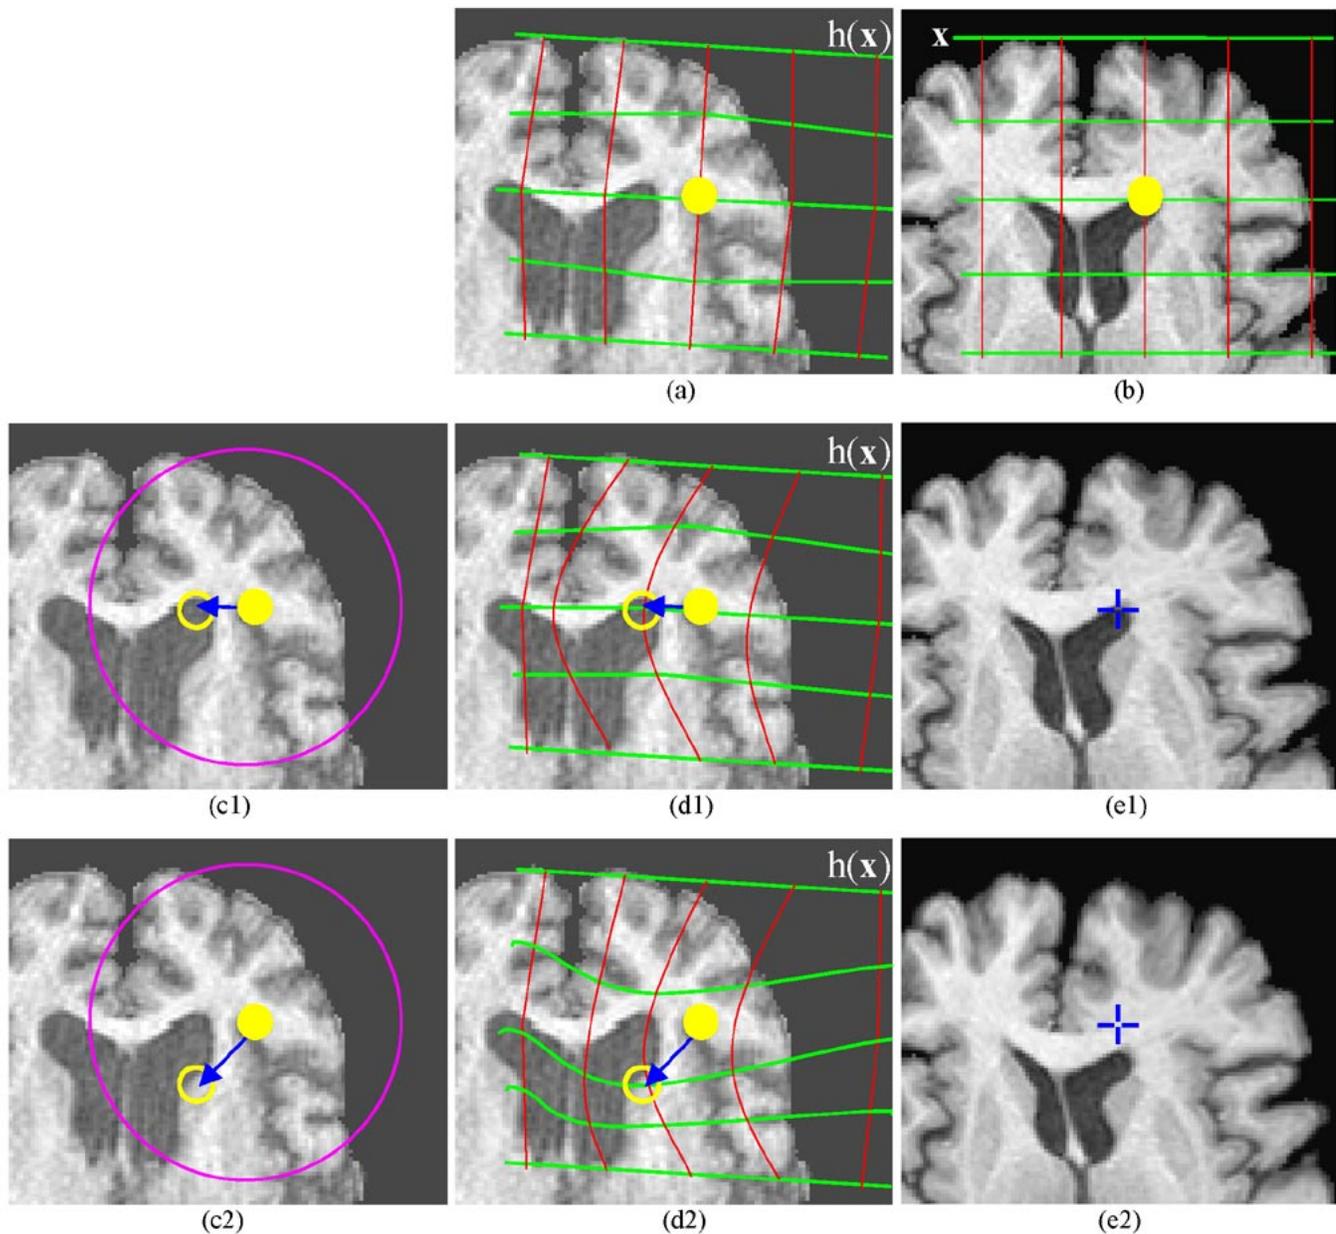


Fig. 4. Demonstration of the subvolume deformation mechanism in HAMMER. The selected parts of the subject and template are respectively displayed in (a) and (b). The regular grid in the template (b) was deformed to be a slightly irregular grid in the subject. All the yellow discs denote the same template voxel under consideration. This voxel searches the subject voxels with similar attribute vectors in its neighborhood [large purple circles in (c1) and (c2)], and finds two nearby subject voxels, as small yellow circles in (c1) and (c2). Then the subvolume of this voxel is tentatively deformed as shown in (e1) and (e2), with the deformed grids $h(x)$ respectively shown in (d1) and (d2). Blue crosses in (e1) and (e2) are at the same positions, which are used to compare the different deformations in these two images.

to several tentative nearby locations, based on which nearby voxels have similar attribute vectors with the voxel under consideration. *It is important to note that the optimal deformation of the subvolume is not determined based on the similarity of one voxel, but is determined by integrating the similarity of the attribute vectors within the whole subvolume.* For example (see Fig. 4), when several tentative deformations of the corner of the ventricle are considered, the similarity of attribute vectors within the whole subvolume around this corner voxel is evaluated. The subvolume of this ventricular corner voxel will be finally deformed to the optimal position, if the total similarity is over a certain threshold (t_{Volume}). Otherwise, this voxel will produce

no deformation on its subvolume. Fig. 4(a) and (b) displays the selected part of the subject and template images. The regular grid x in the template Fig. 4(b) has been deformed slightly to the irregular grid $h(x)$ in the subject Fig. 4(a). The yellow discs in this figure denote the same template voxel under consideration. The search domain of this voxel (a sphere with the radius D_T , which gradually decreases to one with increasing iteration number) is shown as the large purple circles in Fig. 4(c1) and (c2). Two subject voxels with the similar attribute vectors are found in its neighborhood, as small yellow circles in Fig. 4(c1) and (c2). Then the subvolume of this voxel is tentatively deformed to these two positions, as shown in Fig. 4(e1) and (e2). The corresponding

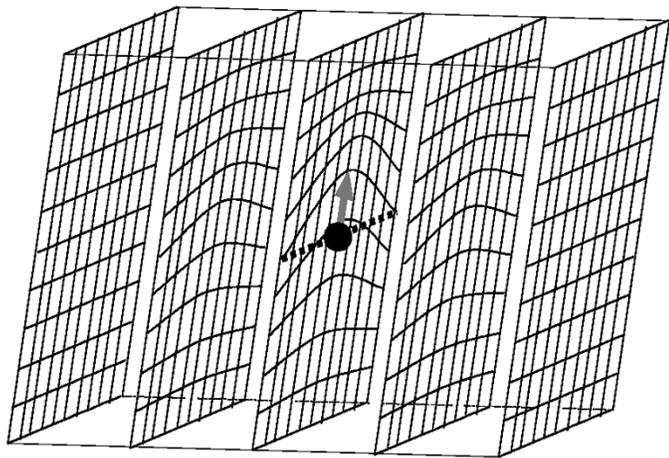


Fig. 5. Schematic representation of the subvolume deformation kernel. Here, the original subvolume is a regular cubic grid, before applying a displacement to the black voxel x . The deformation on the black voxel, at the direction of the gray arrow, brings the deformations to the neighboring voxels, according to a Gaussian falloff kernel, with zero deformations on the boundaries of this subvolume.

tentatively deformed template grids $h(x)$ are given in Fig. 4(d1) and (d2). By integrating the similarity of the template and target attribute vectors in a deformed subvolume, this subvolume is finally deformed to an optimal position such as Fig. 4(e1). Blue crosses in Fig. 4(e1) and (e2) are at the same positions, which are used as a reference in comparing two different deformations in these two images. It should be indicated that the size of the subvolume is defined in a hierarchical fashion. It is initially large, and reduces gradually with time. Also, for speeding up the procedure of evaluating subvolume similarity, the voxels in this subvolume are subsampled according to the size of this subvolume.

The mathematical form for the deformation of the subvolume is described next. Notice that the goal of our deformation mechanism is to minimize the term $\sum_{z \in n(x)} \varepsilon(z)(1 - m(\mathbf{a}_T(z), \mathbf{a}_S(h(z))))$ for each template driving voxel x . When deforming the subvolume of the template driving voxel x , we should produce no discontinuities at the boundary of the subvolume. In order to reach this objective, we propagate the deformation on the template driving voxel x to its neighboring voxel z according to a Gaussian kernel. Note that the current position of the transformed driving voxel x is $h(x)$, and similarly the current position of the neighboring voxel z is $h(z)$ in the subject. Let us assume that the transformation $h(\cdot)$ is updated to map x to a tentative position $h(x) + \Delta$, during the greedy search. Then, the new position of the neighboring voxel z becomes

$$h(z) + \Delta \cdot \exp\left(-\frac{\|x - z\|^2}{2\sigma^2}\right)$$

where σ is a parameter that makes $\exp(-(\|x - z\|^2/2\sigma^2))$ close to zero for the voxels z in the boundary of the subvolume around the template driving voxel x . This definition leaves the boundary voxels unchanged and, hence, maintains the smoothness and continuity of the displacement field. Fig. 5 illustrates the deformation result on an regular cubic subvolume,

by performing the subvolume deformation mechanism on the black voxel x at the direction of the gray arrow. Notably, the deformation along the boundary of this subvolume is zero.

Using this deformation mechanism, Fig. 6 demonstrates some tentative positions of the red-circled voxel and the corresponding deformations of the subvolume around it. Fig. 6(a) is a cross-sectional image of the template, with the green boundary used as a reference for comparing the deformations produced in Fig. 6(b1)–(b4). This template image was tentatively deformed in Fig. 6(b1)–(b4) to the left, top, right, and bottom, as indicated by the yellow arrows. Notably, the deformations displayed in Fig. 6 have fairly global nature, which are particularly applied in the initial deformation stages of HAMMER, while the deformations displayed in Fig. 4 are local and are usually applied in the later deformation stages.

Resolving Ambiguities in Feature Matching: It is important to note that there is residual variability in the cortex, which is primarily due to fundamental morphological differences between the template and the subjects. Typical examples are the “one versus two” sulci, or the absence of a sulcus in a subject, which is present in the template. In such cases, HAMMER is designed to relax the matching forces when no good matches are found, and the deformations on such sulci will be driven primarily by the deformations of the driving voxels in the neighboring structures. In this way, a gyrus or a flat part of the cortex will not be forced to match a sulcus, simply because a sulcus happens to be in its vicinity. In order for a match to be enforced, a high similarity of the attribute vectors must be present. Otherwise, the matching will be driven by other driving voxels, which display a high degree of similarity of their attribute vectors. A typical example is given in Fig. 14. In that example, we can see that a short subject sulcus, indicated by the black arrow, is still kept as a sulcus in the final warping, even though the corresponding part in the template is a gyrus. In cases in which one brain has a single sulcus and another brain has two sulci, HAMMER will find two potential matches. However, none of these matches will seriously influence the deformation, for two reasons. First, it is likely that the attribute vectors in the vicinity of a two-branch sulcus are different from the attribute vectors in the vicinity of a single sulcus. Second, a compromise will be found between forces emanating from the two potentially matching sulci toward the single sulcus, which will further reduce the weight of those forces and the degree of influence on the deformation. Effectively, in such a case, the deformation will be driven primarily by other more definitive matches.

Consistent Transformation Framework: Constructing consistent transformations, i.e., transformations which give identical mapping between two brains regardless of which of the two brains is treated as the template, have gained interest in the medical imaging community during the past two years [42], [43], after the pioneering work in [42]. The concept of consistent transformations has been previously introduced in the computer vision community by [44], for normalization of hand-written characters. In HAMMER, we use the symmetric energy function in (1), which constrains the inverse transformation $h^{-1}(x)$ and pertains to the deformation of “subject to template.” It is quite important to include this second consistency term, particularly for warping the elderly subjects that

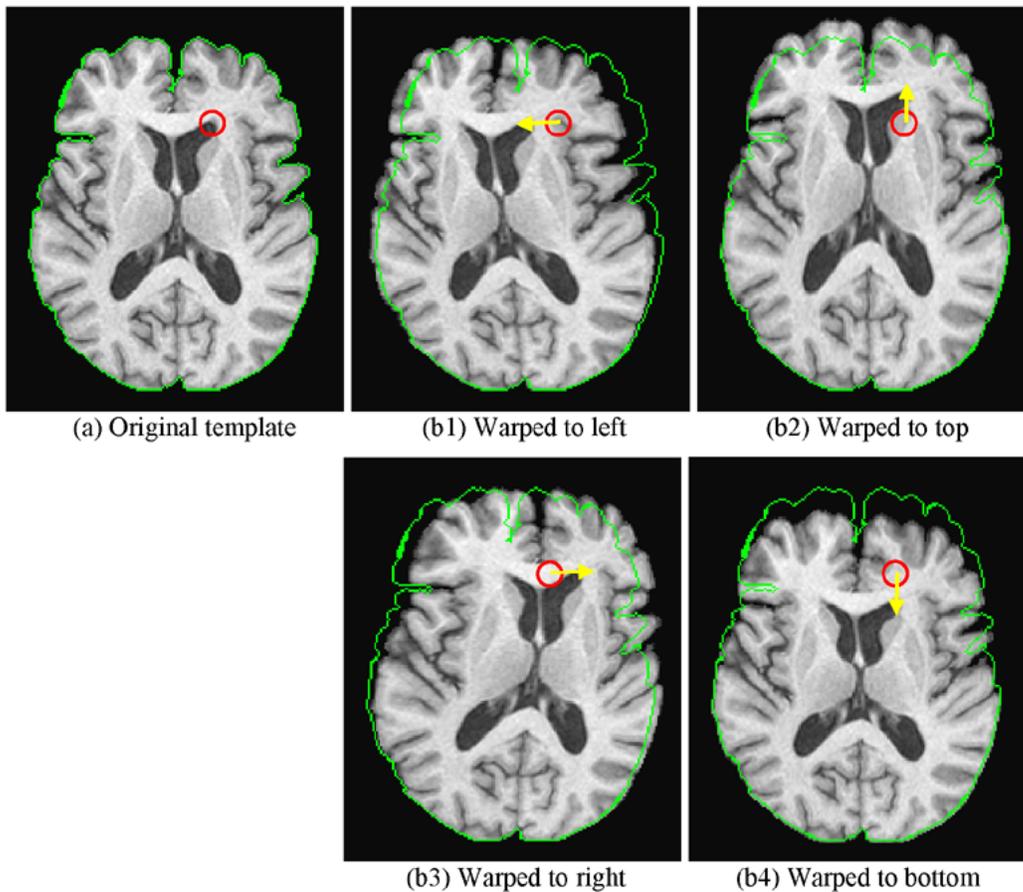


Fig. 6. Demonstrations of tentative deformation of a large subvolume of the driving voxel to the nearby positions. (a) Cross-sectional image of the template. (b1)–(b4) Its four tentative deformations. The red circles are all at the same positions. The green contours in all figures are the same, which are used as the references for observing the respective deformations. Yellow arrows are the displacement vectors of the red-circled voxel, which drives the deformation.

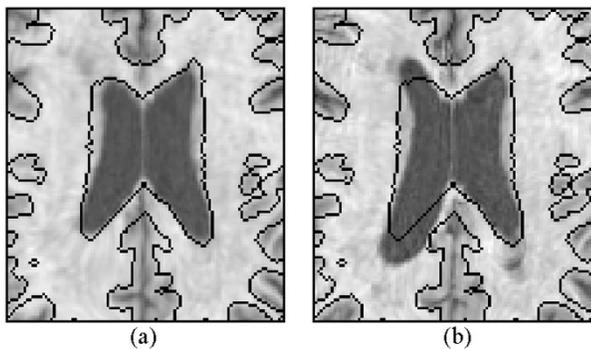


Fig. 7. Demonstration on the importance of including the consistency term in our energy function. (a) Good warping result, obtained using the consistency term. (b) Incorrect warping result of the same subject, obtained without using the consistency term. The black contours, used as reference, are the same WM boundaries of the template brain. The subject used in this figure is the same as the one used in Fig. 9, which has a large ventricle, but the template has an average ventricle.

usually have very large ventricles, to the template that has an average ventricle. This formulation makes our model robust to the initialization and also increases the chances for our model to avoid bad local minima. Fig. 7 vividly demonstrates the importance of including the consistency term into our model, with a good warping result by using the consistency term in

Fig. 7(a), and an incorrect warping result without using the consistency term in Fig. 7(b). The black contours in Fig. 7(a) and (b) are the same WM boundaries of the template brain. These contours are used as references for comparing the warping results.

In the method presented in [42] and [43] for constructing consistent transformations, the forward and reverse transformations are defined separately, and solutions that make these transformations consistent are favored. In this paper, we only define one transformation, $h(x)$, from the template to the subject, and implicitly enforce consistency by using $h^{-1}(x)$ for the reverse mapping. Eventually, we solve for a single transformation. We know that the first energy term of (1) is defined in the template space and used to calculate the transformation $h(x)$. The second energy term is defined in the subject space and used to constrain the inverse transformation of $h(x)$, $h^{-1}(x)$. In the beginning of this section, we provided a hierarchical deformation method for determining the transformation $h(x)$ in order to match the template with the subject. We next describe a method for converting the requirements from the inverse transformation $h^{-1}(x)$ to an additional constraint on the transformation $h(x)$. For example, according to the first energy term, a template driving voxel F is deformed to a subject voxel P by the transformation $h(x)$. But according to the second energy term, this template driving

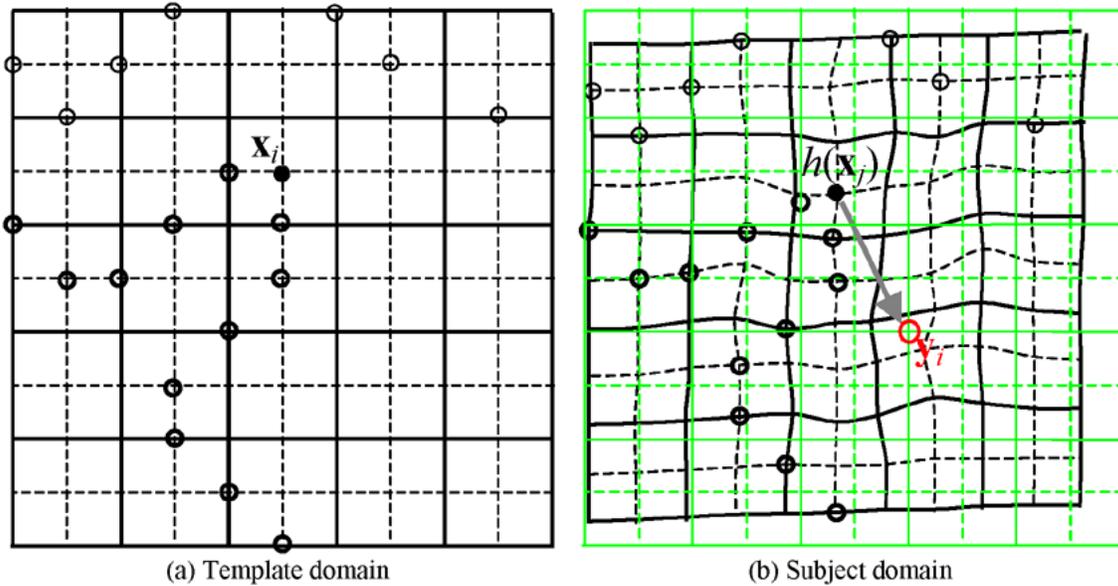


Fig. 8. Schematic description of our strategy for converting the requirements from the inverse transformation $h^{-1}(x)$ to an additional constraint on the transformation $h(x)$. (a) Regular black lattice in the template domain before displacement, where the black circles denote the template driving voxels. The regular black lattice with the driving voxels was deformed to the black irregular lattice in the subject domain in (b). The green regular lattices in (b) represent the locations of the subject voxels. For the subject driving voxel y_i , in its neighborhood we find a displaced template driving voxel $h(x_j)$ with the most similar attribute vector. If the degree of the similarity of the attribute vectors between the two voxels x_j and y_i is over a certain threshold, then a force that is displayed as a gray arrow is produced on $h(x_j)$, for driving the template subvolume around $h(x_j)$ to the subject driving voxel y_i .

voxel F should correspond to a subject voxel Q . This implies that the inverse transformation $h^{-1}(x)$ will require that the subject's voxel Q deform to the template's driving voxel F . Accordingly, we simply produce a force on this template driving voxel F at the direction of FQ , which updates the transformation $h(x)$ in the subvolume of the template voxel F . To speed up our algorithm, we enforce consistency only on the subject driving voxels, since these voxels effectively define correspondence between the template and the subject. Consistency in the remaining, nondriving voxels can be achieved by designing consistent interpolants, i.e., interpolants that consistently propagate correspondences from the driving voxels to the rest of the image volume. Designing such interpolants is beyond the scope of this paper.

Let us assume that, in the current iteration, the set of the template driving voxels is $\{x_i | 1 \leq i \leq N_T\}$, such as a number of small black circles in Fig. 8(a). Similarly, let $\{y_i | 1 \leq i \leq N_S\}$ be the set of the N_S subject driving voxels. [This latter set of driving voxels is fixed in the whole deformation procedure; Fig. 3(b) shows an example of the subject driving voxels]. With the transformation $h(x)$ that is calculated in the previous iteration, we know the set of the template driving voxels $\{x_i | 1 \leq i \leq N_T\}$ has been displaced into $\{h(x_i) | 1 \leq i \leq N_T\}$, such as the small black circles in Fig. 8(b). Then for each subject driving voxel y_i , such as the red point in Fig. 8(b), in its neighborhood (a sphere with the radius D_S) we search for a displaced template driving voxel $h(x_j)$, where $h(x_j) \in \{h(x_i) | 1 \leq i \leq N_T\}$, with the most similar attribute vector. If the degree of the similarity between the subject driving voxel y_i and the template driving voxel x_j is over a certain similarity threshold t_{voxels} , then a force is applied to the displaced template voxel $h(x_j)$ at the di-

rection from $h(x_j)$ to y_i , as shown as a gray arrow in Fig. 8(b). We used 0.8 for the initial similarity threshold, which decreases gradually to 0.01 with increasing number of iterations. Finally, the displaced template driving voxel $h(x_j)$, including its subvolume, will be deformed according to the subvolume deformation mechanism.

Smoothing of Displacement Fields: In the end of each deformation iteration, we use two smoothing techniques to refine the displacement fields. The first technique uses a global affine transformation to constrain the global consistency of the displacement fields; similarly, the locally calculated affine transformations can also be used to constrain the local consistency of the displacement fields. This is particularly important in the initial deformation stages, when the driving voxels are sparse in the image space and thereby more easily affected by the noise. The global affine transformation can be simply calculated from the positions of the template driving voxels before and after the current iteration. Let us assume that the transformation fields on the set of the template driving voxels $\{x_i | 1 \leq i \leq N_T\}$ are $\{h(x_i) | 1 \leq i \leq N_T\}$ before the current iteration, and $\{h(x_i) + \Delta_i | 1 \leq i \leq N_T\}$ after the current iteration. Here, Δ_i is a current deformation on the template driving voxel x_i , due to the current iteration. Then, the global affine transformation matrix is obtained by using a least squares method to solve the variables A and B of

$$\begin{aligned} [h(x_1) + \Delta_1, \dots, h(x_i) + \Delta_i, \dots, h(x_{N_T}) + \Delta_{N_T}] \\ = A[h(x_1), \dots, h(x_i), \dots, h(x_{N_T})] + B \end{aligned}$$

where $h(x_i)$ is a three-element column vector, A is a 3×3 transformation matrix, and B , as a shift factor, is a three-element

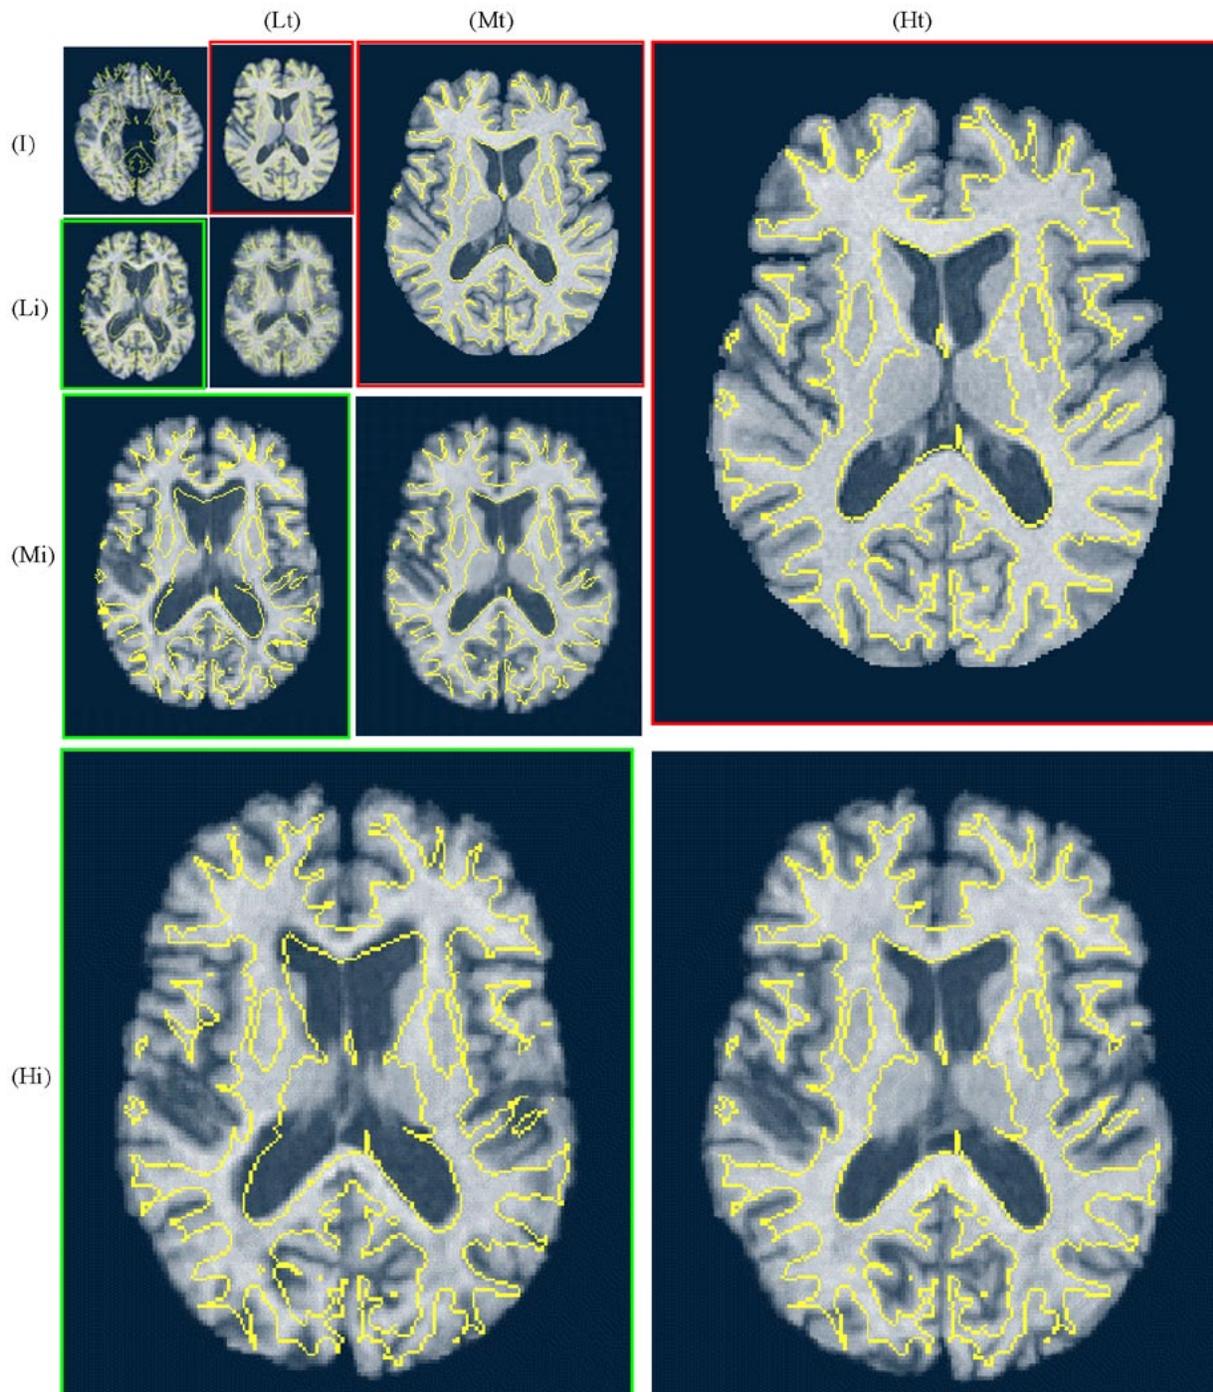


Fig. 9. Example of warping a subject brain image to the template. The bottom right panel shows the final result. All yellow contours are the same, although displayed at different sizes for different resolution levels. These yellow contours are the WM boundaries of the template, and are used here as a reference. The four small images in the upper left corner are examples in the low resolution. They are, respectively, the original subject image (I), the globally affine-transformed version of the original subject that is used as the initialization (Li), the template image (Lt), and the final warping subject in the low resolution. The image in (Mi) is the initial warping subject in the middle resolution, using the displacement fields that were upsampled from the low resolution. The image in (Mt) is the template in the middle resolution. The final warping image in this middle resolution is put in the middle of this figure. For the high-resolution procedure, three slice images are shown. The image in (Hi) is the initial warping result that was upsampled from the middle resolution, the image in (Ht) is the template, and the image in the bottom right panel is the final warping result of the subject, overlaid on the template's WM boundary.

column vector. The transformation field on each template voxel is then refined as

$$\lambda \cdot (h(x) + \Delta) + (1 - \lambda) \cdot (A \cdot h(x) + B)$$

where Δ is the deformation of the voxel x from the current iteration, and λ is a parameter ranging from zero to one. The parameter λ is chosen close to 0 in the initial deformation stage (i.e., more global constraints) and it increases gradually to one

(i.e., less global constraints) with the increase of the iteration number. That is, the global constraints are dominant in the initial deformation stages, which helps the algorithm avoid local minima.

The second smoothing technique is the use of the third smoothness constraint in (1). The Gaussian kernel that we selected for deforming the subvolume does not make the Laplacian term close to zero everywhere. The reason that we selected the Gaussian function as the kernel is because the Gaussian function is a simple smooth function and it prescribes that the influence of the driving voxel on the displacement of other points fades away with distance. Also, according to our knowledge, there is no kernel that can meet our requirements which also makes the Laplacian term zero everywhere. In this way, we jointly apply this Laplacian smoothness constraint in deforming the subvolume of each driving voxel. Moreover, we smooth again the displacement fields in the whole volume in the end of each iteration, according to the constraint of the Laplacian cost term. This leads to displacement fields well behaved, in that they are smooth and have smooth first derivatives. This can be observed from the displacement fields in Fig. 10.

Multiresolution Formulation: In order to speed up HAMMER, reduce local minima, and create smooth deformation fields, we have implemented our algorithm in a multiresolution framework. Specifically, we used three different levels, corresponding to the original image resolution (high-resolution), a subsampled version by a factor of two (mid-resolution), and a subsampled version by a factor of four (low-resolution). For each resolution, we first calculate the attribute vector for each voxel in the template and subject. Then we use HAMMER to register them and warp the subject to the template. The transformation function $h(x)$ in one low resolution are upsampled and linearly interpolated to the next resolution. The upsampled transformation function is then used to generate the initialization for the current resolution.

In Fig. 9, we give an example of warping a subject brain image to the template using this multiresolution formulation. The explanation for all subfigures in Fig. 9 is provided in Section III.

D. Summary of the HAMMER Algorithm

There are two major stages in our brain warping algorithm. The first stage is to remove the skull and label tissues as GM, WM, and CSF based on the intensities, as described in [39]. The second stage is to use the fully automatic HAMMER algorithm to register and warp the brain images. The detailed procedure of HAMMER is described next. This procedure is repeatedly used from low resolution to high resolution.

1) Transform the subject to the space of the template using a global affine transformation that is determined by a matrix of the regular moments of the subject [34].

2) Compute an attribute vector for each voxel in the template and the subject. Since we have labeled tissues as GM, WM, and CSF, we can calculate the respective attributes from the local structures of GM, WM, and CSF for each voxel. Notice that

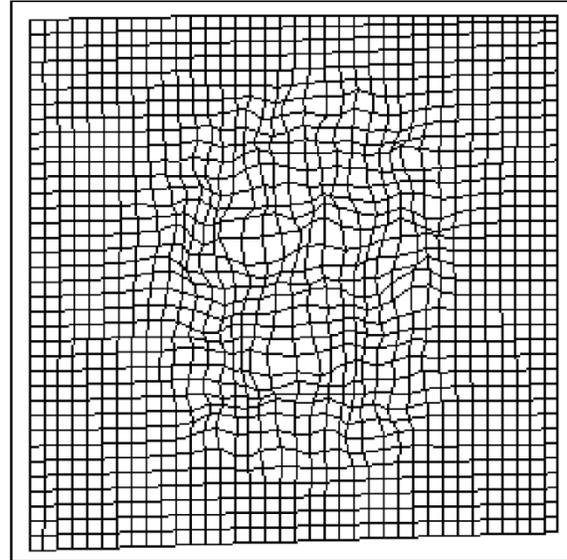


Fig. 10. Typical displacement field for the brain warping example shown in Fig. 9.

there are 13 GMIs that can be calculated from each tissue. In the experiments of the next section, we used the first three GMIs in order to speed up performance, and because we found those to be adequate for those experiments.

3) Determine a set of the driving boundary voxels $\{y_i | 1 \leq i \leq N_S\}$ in the subject, based on their attribute vectors. This set of subject driving voxels is fixed during the whole warping procedure in this paper.

4) Hierarchically select a set of the boundary voxels in the template, $\{x_i | 1 \leq i \leq N_T\}$, for driving the deformation of the template. The initial set of the template driving voxels includes only the distinctive template voxels, such as roots of sulci and crowns of gyri. In the later iterations, other boundary voxels such as ventricular boundaries are added to this set. In the final stage, all voxels in the template are included in this voxel set.

5) For each subject driving voxel y_i , search in its neighborhood (a sphere with the radius D_S) for a displaced template driving voxel $h(x_j)$ with the most similar attribute vector. If the degree of similarity between the subject driving voxel y_i and the template driving voxel x_j is above a threshold t_{Voxels} (t_{Voxels} is initially high, and decreases gradually with time), then a force is produced on the deformed template driving voxel $h(x_j)$, in the direction from $h(x_j)$ to y_i .

6) For each template driving voxel x_i , search in its neighborhood (a sphere with the radius D_T) for all subject voxels with similar attribute vectors. (The required similarity is initially high, and decreases gradually as the model converges.) Then, tentatively deform the subvolume of the voxel x_i to each similar subject voxel that has been found, and integrate the similarity degrees of all voxels in the subvolume. This subvolume will be finally deformed to a subject voxel with the largest similarity degree, if this largest similarity degree is over a certain threshold (t_{Volume}). Additionally, if there exist forces performed on the template driving voxel x_i from the subject driving voxels,

TABLE I
LIST OF THE PARAMETERS USED IN THIS PAPER. MOST OF THEM EVOLVE AS THE ALGORITHM PROGRESSES.
WE USED THE SAME SET OF PARAMETERS FOR THE ALL RESULTS SHOWN IN THIS PAPER.

Parameters	Low level	Middle level	High level
<i>Iter</i> : The maximum iteration number.	50		
δ : A general search range.	12	10	8
<i>i</i> : The current iteration number.			
$\tau = \frac{i}{Iter}$			
β : A control parameter in equation 1.	0.5		
<i>R</i> : The radius of the spherical subvolume, used to calculate GMIs.	3	3	7
$n(\mathbf{x})$: The spherical subvolume, with the radius <i>r</i> .	$r = 0.5\delta e^{\left(\frac{\tau^2}{2.0 \times 0.16}\right)} + 1$		
D_T : The radius of the spherical search domain, used for the template's driving voxel.	$D_T = r$		
D_S : The radius of the spherical search domain, used for the subject's driving voxel.	$D_T = r + 6$		
t_{Voxels} : The threshold for the similarity of the attribute vectors in the subject and the template, respectively.	$t_{\text{Voxels}} = 0.8(1 - \tau) + 0.001$		
t_{Volume} : The threshold for the similarity of attribute vectors in two subvolumes.	$t_{\text{Volume}} = t_{\text{Voxels}}$		
σ : The variance of the interpolating Gaussian kernel.	$\sigma = \frac{r}{3.0}$		
λ : A parameter used to control the global constraint.	$\lambda = 0.25 + 0.75e^{\left(\frac{(\tau-1)^2}{2.0 \times 0.25^2}\right)}$		

which are determined in Step 5, then this subvolume will also be deformed under these forces.

7) Hierarchically refine the displacement fields using local and global affine transformations that are calculated from the deformations in the template driving voxels.

8) Refine the displacement fields using the smoothness constraint as described as the third energy term in (1).

9) If the warping procedure has converged, then stop. Otherwise, go back to Step 4.

The registration accuracy is definitely related to the selections of the parameters that are included in our algorithm. Moreover, some parameters are highly correlated, that is, once one parameter is changed, other parameters should be changed accordingly. For example, the registration accuracy is related to the spatial distribution of driving voxels, the variance of the Gaussian kernel, and also the stopping criteria of our algorithm. We require that the driving voxels be somewhat uniformly distributed in the brain space; otherwise, the deformation in one part of brain will probably dominate over other parts. This is particularly important for the initial deformation stages, where the number of the driving voxels is relatively small. The size of the interpolating Gaussian kernel is related to the number of the driving voxels. It is initially large, since the number of the driving voxels is small at those stages. It then decreases gradually as the number of the driving voxels increases. Notably, most parameters that are used in our algorithm are designed to evolve as the algorithm progresses. Therefore, the algorithm is ensured to converge in the end. We can simply stop the algorithm, when the deformation in the current iteration is less than one amount.

In our experiments, we used a maximum iteration number of 50 for each resolution.

Table I provides our choice of the parameters. Most parameters evolve as the algorithm progresses. Also, most parameters are related; given one, other ones can be obtained. The basic parameters are the general search size δ and the current iteration order *i*.

III. RESULTS

We describe three different experiments to demonstrate the performance of HAMMER. The first experiment demonstrates the procedure for hierarchically warping a subject brain to a template brain, where the subject was selected from a group of elderly individuals with high degree of atrophy. The second experiment demonstrates the excellent performance of our algorithm on data from 18 subjects, with a very clear average brain image after warping to the same template. The third experiment provides a validation of our algorithm using a set of simulated deformations. Without any code optimization, our algorithm currently requires about two hours of processing on an SGI OCTANE workstation (300 MHz) for a full 3-D warping of two magnetic resonance (MR) images. Notice that we used the same parameters for all results in this paper.

A. Single-Subject Experiment

This experiment provides a procedure of hierarchically warping a subject to the template, as shown in Fig. 9. The subject used in this experiment is an elderly individual with significant atrophy and ventricular enlargement. Three different

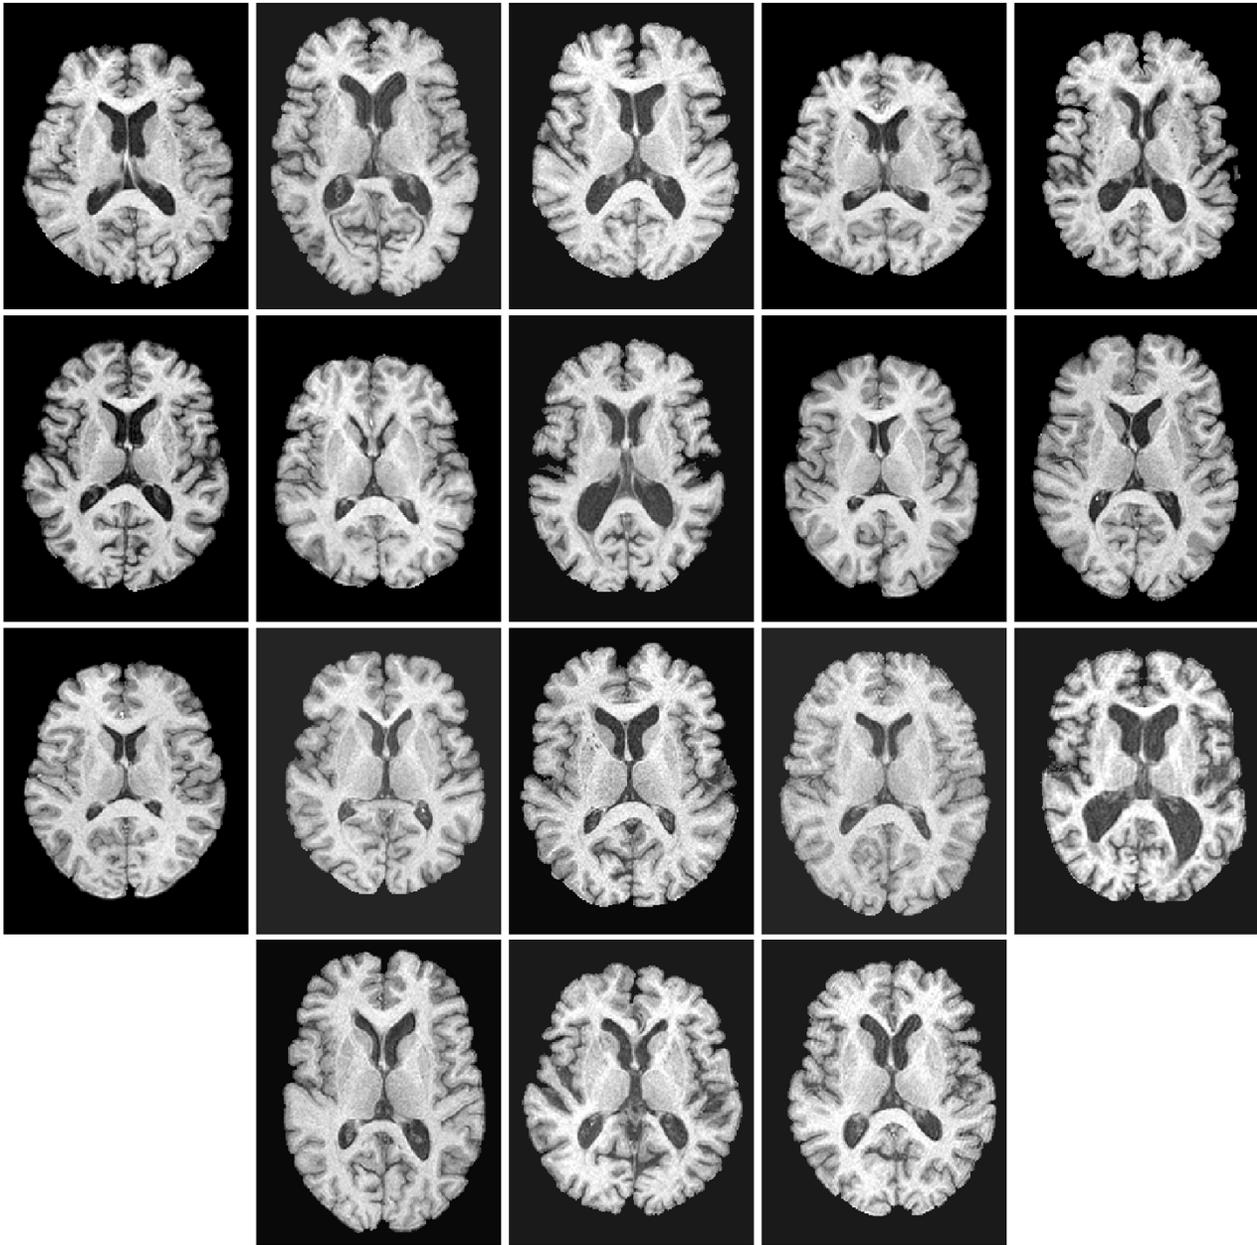


Fig. 11. Eighteen randomly selected subjects, which are used to construct the average brain image.

resolution levels are used, i.e., low, middle, and high resolution. For each resolution level, the cross-sectional images in Fig. 9 are taken from the same slice. The yellow contours are taken from the same volume and delineate the WM boundary of the template; we use them as a reference for comparison. The small images displayed in the upper left corner are taken from the low-resolution warping procedure. In the beginning of the warping procedure, the original subject image in Fig. 9(I) was automatically transformed to the space of the template in Fig. 9(Lt) using a global affine transformation [34]. The effect of this affine transformation can be observed by comparing the three images in Fig. 9(I), (Li), and (Lt); notice that the image in Fig. 9(Li) is the affine-transformed version. The final warping result in the low resolution is displayed under the template image in Fig. 9(Lt). The corresponding displacement field was

upsampled and used for initialization at the middle resolution, as displayed in Fig. 9(Mi). This initial warping result was improved in the end of the middle resolution, by comparing the template image in Fig. 9(Mt) with the warping result under this template. Similarly, in the beginning of high resolution stage, the displacement fields were upsampled and interpolated to the high resolution and, therefore, we obtained the initial warping result in this high resolution in Fig. 9(Hi). The final warping result of the subject is displayed in the right corner. A comparison of this final warping result with the template in Fig. 9(Ht) indicates the accuracy of our algorithm. Fig. 10 shows a typical displacement field obtained using HAMMER. We can observe the relative smoothness of the associated displacement fields, despite the initial notable differences such as in the ventricles of the subject and the template.

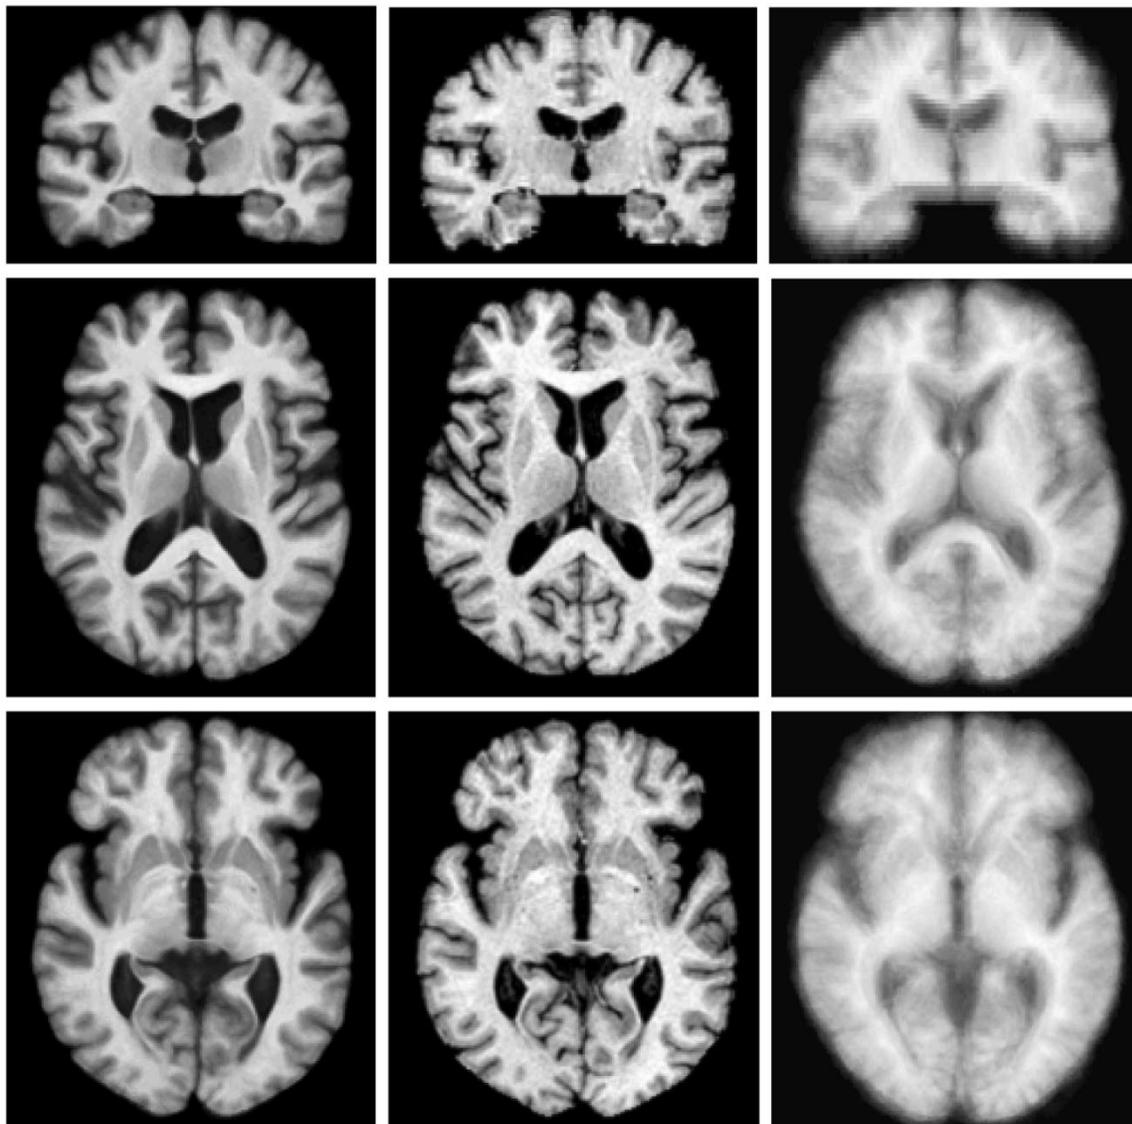


Fig. 12. Average brain image (left) reconstructed from 18 very different subjects in Fig. 11. The template images in the same slices are also displayed for comparison. The sharpness of the average is indicative of the registration accuracy obtained by HAMMER.

B. Experiment on Constructing Average Brain Image

The sharpness of the average of the warped brain images is often used as a visual display of the accuracy of the warping algorithm. We selected 18 subjects (shown in Fig. 11), with a wide variability of ventricular shape and size, and, of course, representative of variability in other structures. We used HAMMER to fully automatically warp these 18 subjects to the template, using the same default set of parameters. Fig. 12 shows the average of those 18 warped brain images, in three different views. For comparison, the same cross-sectional images in the template are also displayed. Notably, the cortical regions, ventricles, and WM/GM boundaries are very clear. Moreover, the caudate nucleus and the lenticular nucleus are also very clear. That indicates the high-registration accuracy obtained by HAMMER.

For comparing the similarity between the average brain and the template brain, we also generated their 3-D renderings, as shown in Fig. 13. They are very similar. Particularly, a very

clear definition of the cortical anatomy is evident in this average brain, and it is in agreement with the template image. For better displaying the deformation in the cortex, we selected a typical brain and displayed its 3-D renderings before and after HAMMER-based warping (Fig. 14). Note that the cortex of the warped subject is very close to that of the template as shown in Fig. 13. It is also very important to note that the characteristic details of the cortical anatomy were maintained through this warping procedure. Particularly, the short sulcus, indicated by the black arrows in Fig. 14, was maintained as a sulcus, even though its corresponding location in the template corresponds to a gyrus. This is achieved by the strategy of resolving ambiguities in feature matching given in Section II-C.

C. Validation

In addition to the registration accuracy revealed by the 18-brain average image of the previous experiment, we

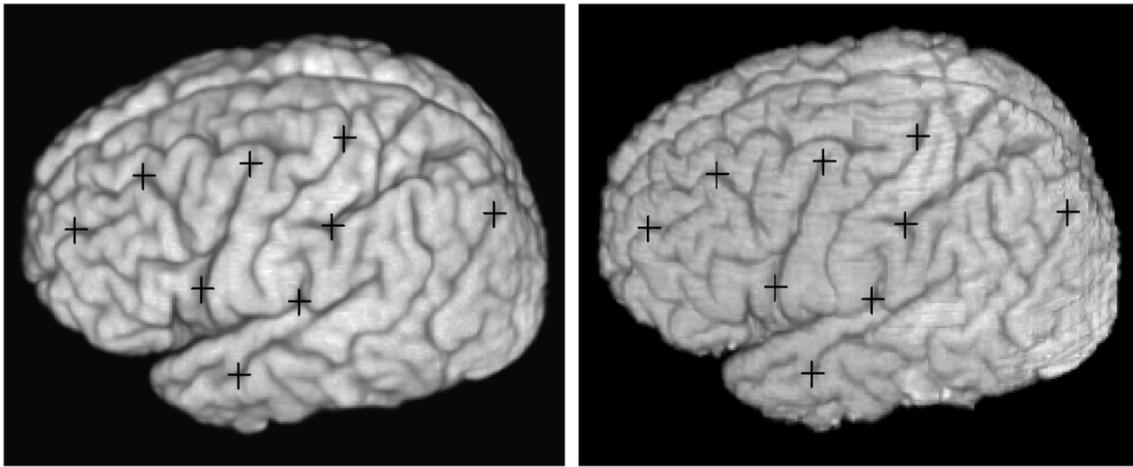


Fig. 13. The 3-D renderings of (a) the average brain and (b) the template brain. The black crosses, with the same 3-D coordinates respectively in (a) and (b), are used as references for comparing the correspondences in the average and template brains.

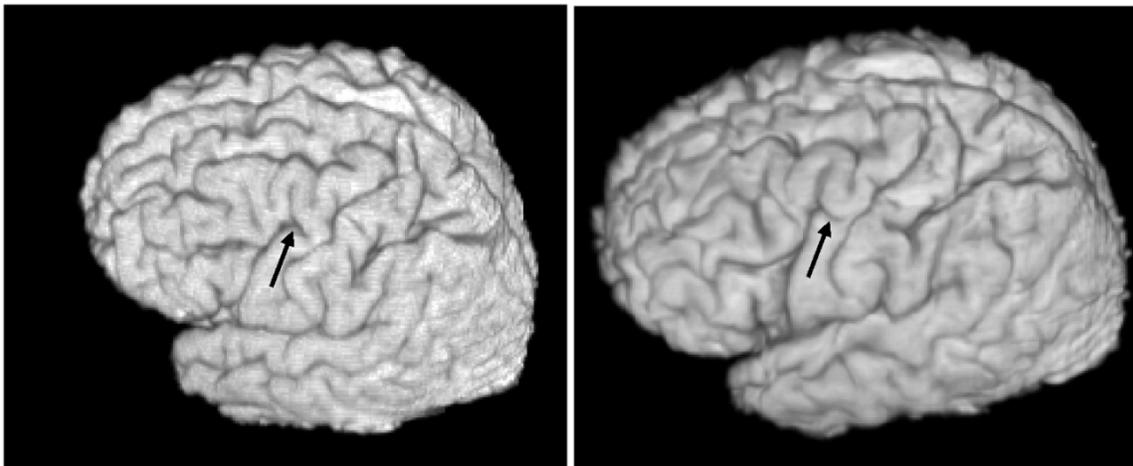


Fig. 14. Comparison on one subject's original image with its warped configuration, using the template of Fig. 13(b). The black arrows indicate the same short sulcus before and after HAMMER-based warping. This shows that a sulcus will not be forced to match a gyrus, simply because a gyrus happens to be in its vicinity.

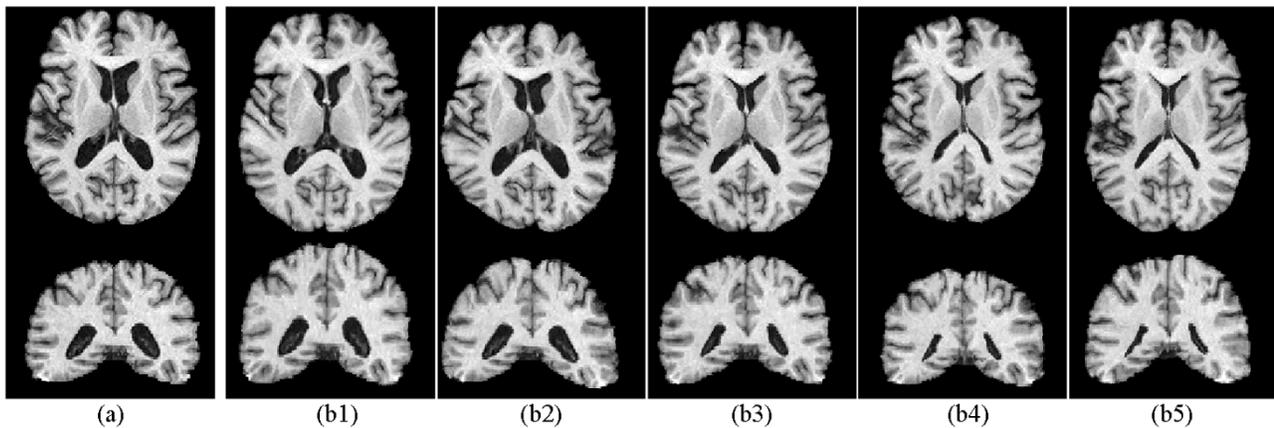


Fig. 15. Synthesized brain images used to validate the registration performance of HAMMER. (a) Selected subject. (b1)–(b5) Five deformed versions of the subject in (a), where different degrees of atrophy were simulated by introducing pressure within the ventricles, and changes in the cortical folding pattern were achieved by manually drawing a number of sulcal curves along the outer cortex in five real brains, and using them as constraints in driving a 3-D elastic warping deformation.

used simulated deformations to validate the performance of HAMMER. The simulation data was created as follows. First, we randomly selected a brain, which is shown in Fig. 15(a).

We then manually painted the precentral gyrus and the superior temporal gyrus, as shown in Fig. 16. We need to note, here, that any region of interest could be selected for this validation

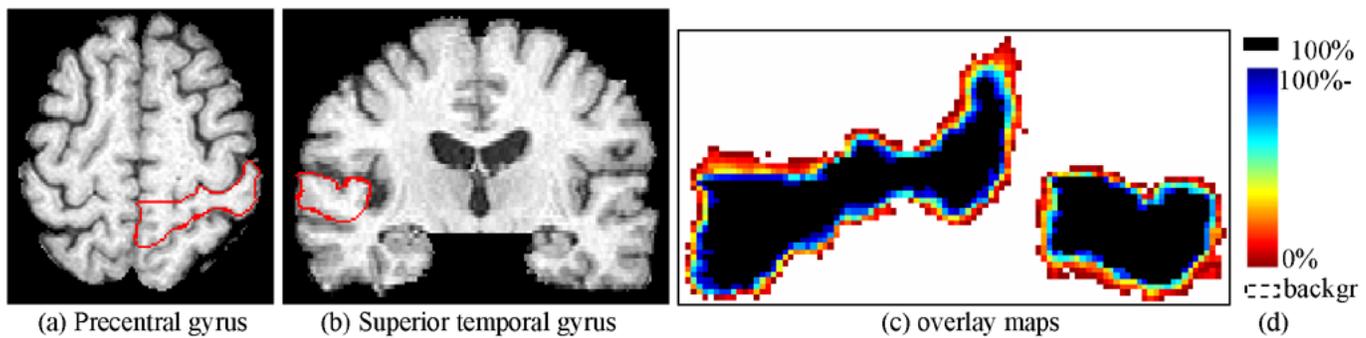


Fig. 16. Two structures used for the validation of HAMMER. (a) Precentral gyrus. (b) Superior temporal gyrus. (c) Color codings of the overlay maps of precentral gyrus and superior temporal gyrus from five restored brains, with 100% overlay as black and the background as white. (d) Color coding bar, with the overlay percentages.

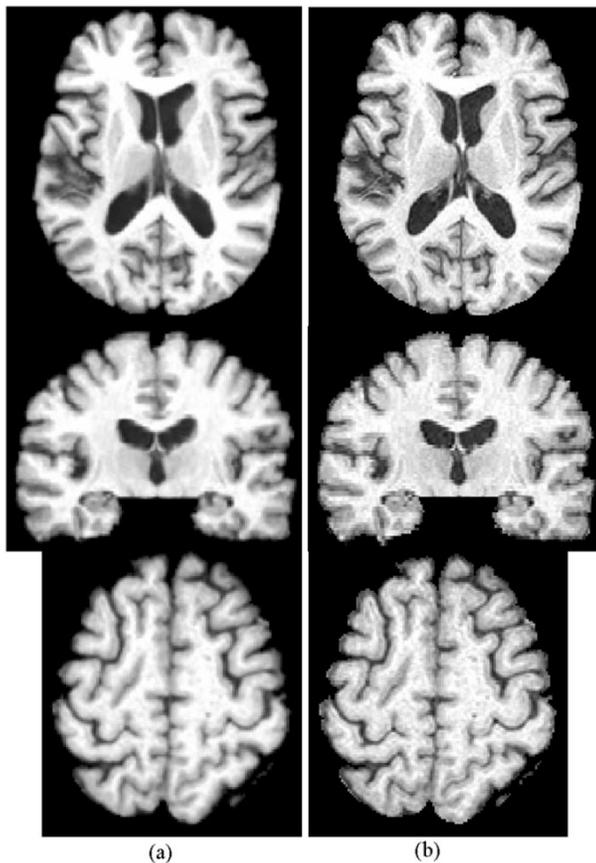


Fig. 17. Average brain constructed from five synthesized brain images, after warping to the same template. This is used to visually examine the registration accuracy of our HAMMER algorithm.

experiment, which intended to demonstrate that HAMMER is able to retrieve simulated deformations. We selected these two regions of interest (ROIs) simply because they have an anatomical meaning. We then use the STAR elastic warping method [39] to warp this selected brain by placing forces on the ventricles and the cortex, thereby obtaining five synthesized brains, as shown in Fig. 15(b1–b5). Notably, this warping procedure was based on a number of sulcal constraints placed on the outer cortex, which were obtained from five real brain images in our database. We selected five real brain images, so that the resulting simulated deformations correspond to

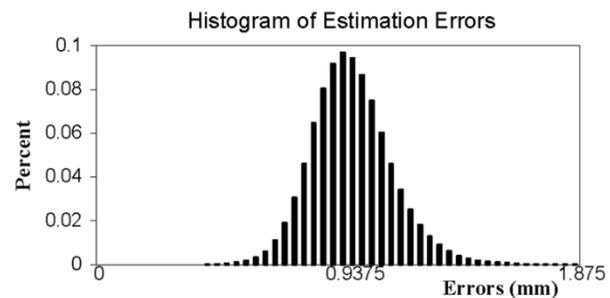


Fig. 18. Histogram of the estimation errors between the simulated deformations and the HAMMER-estimated deformations. The average error is 0.966 mm, and the maximal error 1.875 mm. Notably, there are only 0.01% voxels that have this maximal error 1.875 mm, i.e., only 104 voxels in the volumetric image.

realistic looking brains, which would not be the case if, for example, we randomly perturbed a number of landmarks. The manually painted structures were warped by the same deformation fields. Finally, we used HAMMER to *completely independently* warp these five synthesized brain images back to the original space of the selected brain, and thereby obtained five restored brains. A perfect overlap of the painted ROIs implies a perfect registration accuracy. The average overlap of the painted regions was 89.5%, and the average volume difference was 4.0%. Since the structures we selected are very small, much of this error was probably due to small discrepancies along the boundary, most of which are caused by the interpolation that took place in the forward and the backward deformation of the ROIs. This is shown visually in Fig. 16(c), which displays overlay maps, where white indicates 100% overlay. Furthermore, we computed the average brain image from the five restored brain images, which can also be used to visually examine the registration performance of our algorithm. The average brain is displayed in Fig. 17(a), whose comparison with the original brain shown again in Fig. 17(b) indicates the method's accuracy in this simulation experiment. Finally, we quantitatively compared the deformation values voxel-by-voxel in the image, between the simulated deformations and the HAMMER-estimated deformations. Fig. 18 shows the distribution of the measured displacement field error, with the average error 0.966 mm and the maximal error 1.875 mm. Notably, there are only 0.01% voxels that have this maximal error 1.875 mm, i.e., only 104 voxels in the volumetric image. Since the

average error is less than the cubic root of the voxel's volume in these images, we can conclude that *the method achieved subvoxel registration accuracy in this experiment.*

IV. CONCLUSION

We have presented a new approach to deformable registration of brain images, which has resulted in very high accuracy in superposition of images from different subjects. The proposed approach has two major novelties. First, it utilizes an attribute vector, i.e., a number of attributes that are derived from volumetric images, and which reflect geometric characteristics of the underlying anatomical structures. Part of the attribute vector is a set of GMI's whose values depend on the underlying anatomy. The attribute vectors are essential in establishing anatomical correspondences in the deformation procedure, to the extent that anatomical correspondences are reflected by shape similarity. This is a fundamental deviation of our method from other volumetric deformation methods, which are typically based on maximizing image similarity. The second characteristic of our approach is its hierarchical deformation mechanism, which renders it robust to suboptimal solutions. In particular, parts of the anatomy that have distinct shape properties are used as "anchors" during initial stages of the deformation procedure, while the remaining parts simply follow the deformation. As a deformable template deforms toward its target configuration, other parts of the anatomy gradually affect the deformation, as those get close to their respective targets and thus avoid potentially spurious targets. Moreover, the deformation mechanism deforms subvolumes, i.e., relatively large parts of the image at a time, rather than individual voxels, and evaluates the similarity of attribute vectors over the entire subvolume. This adds realism to the deformation mechanism, as, for example, the similarity of a gyrus as a whole would be evaluated, relative to its neighborhood, rather than evaluating similarity on individual voxels.

We have evaluated this algorithm on images from an aging study, which display several difficulties, including low tissue contrast, significant atrophy, and motion artifacts. The very clear averages after HAMMER-based warping (see Figs. 12 and 13), the relative smoothness of the associated displacement fields (see Fig. 10), and our quantitative error measurements that showed subvoxel registration accuracy in simulated deformation experiments indicate that HAMMER will allow us to examine morphology at a very detailed level.

Future work includes a deeper understanding of the relationship between attribute vectors and anatomical features, and on speeding up performance of the algorithm.

REFERENCES

- [1] K. J. Friston, A. P. Holmes, K. J. Worsley, J. P. Poline, C. D. Frith, and R. S. J. Frackowiak, "Statistical parametric maps in functional imaging: A general linear approach," *Human Brain Mapping*, pp. 189–210, 1995.
- [2] J. C. Gee, M. Reivich, and R. Bajcsy, "Elastically deforming 3-D atlas to match anatomical brain images," *J. Comput. Assist. Tomogr.*, vol. 17, pp. 225–236, 1993.
- [3] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "3-D brain mapping using a deformable neuroanatomy," *Phys. Med. Biol.*, vol. 39, pp. 609–618, 1994.
- [4] C. Davatzikos, "Mapping of image data to stereotaxic spaces: Applications to brain mapping," *Human Brain Mapping*, vol. 6, pp. 334–338, 1998.
- [5] C. Davatzikos, A. Genc, D. Xu, and S. M. Resnick, "Voxel-based morphometry using the RAVENS maps: Methods and validation using simulated longitudinal atrophy," *Neuroimage*, vol. 14, pp. 1361–1369, Dec. 1, 2001.
- [6] P. Thompson and A. W. Toga, "A surface-based technique for warping three-dimensional images of the brain," *IEEE Trans. Med. Imag.*, vol. 15, pp. 402–417, Aug. 1996.
- [7] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans, "Automatic 3-D intersubject registration of MR volumetric data in standardized Talairach space," *J. Comput. Assist. Tomogr.*, vol. 18, pp. 192–205, 1994.
- [8] F. L. Bookstein, "Thin-plate splines and the atlas problem for biomedical images," in *Proc. Int. Conf. Information Processing in Medical Imaging*, 1991, pp. 326–342.
- [9] P. A. Freeborough and N. C. Fox, "Modeling brain deformations in Alzheimer's disease by fluid registration of serial MR images," *J. Comput. Assist. Tomogr.*, vol. 22, pp. 838–843, 1998.
- [10] G. Subsol, J. P. Thirion, and N. Ayache, "Application of an automatically built 3-D morphometric brain atlas: Study of cerebral ventricle shape," in *Visualization in Biomedical Computing, Lecture Notes in Computer Science*. Hamburg, Germany: Springer-Verlag, 1996, vol. 1131, pp. 373–382.
- [11] J. Ashburner, C. Hutton, R. Frackowiak, I. Johnsrude, C. Price, and K. Friston, "Identifying global anatomical differences: Deformation-based morphometry," *Human Brain Mapping*, vol. 6, pp. 348–357, 1998.
- [12] V. Megalooikonomou, E. H. Herskovits, and C. Davatzikos, "Computer simulation for evaluation of statistical methods for detection of lesion-deficit associations," *Human Brain Mapping*, vol. 10, no. 2, pp. 61–73, 2000.
- [13] M. Vaillant, C. Davatzikos, R. H. Taylor, and R. N. Bryan, "A path-planning algorithm for image guided neurosurgery," in *Proc. CVRMed II—MRCAS III*, Mar. 1997, pp. 467–476.
- [14] S. K. Kyriacou, C. A. Davatzikos, S. J. Zinreich, and R. N. Bryan, "Non-linear elastic registration of brain images with tumor pathology using a biomechanical model," *IEEE Trans. Med. Imag.*, vol. 18, pp. 580–592, July 1999.
- [15] R. N. Bryan, C. Davatzikos, M. Vaillant, J. L. Prince, S. Letovsky, R. Raghavan, W. Nowinski, G. Salamon, N. Murayama, O. Levrier, and M. Zilbovicius, "Creation of population-based anatomic atlases with a brain image database," in *Proc. 1st Int. Conf. Functional Brain Mapping*, 1995, p. 72.
- [16] T. Peters, B. Davey, P. Munger, R. Comeau, A. Evans, and A. Olivier, "Three-dimensional multimodal image guidance for neurosurgery," *IEEE Trans. Med. Imag.*, vol. 15, pp. 121–128, Apr. 1996.
- [17] F. L. Bookstein, "Principal warps: Thin-plate splines and the decomposition of deformations," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, pp. 567–585, June 1989.
- [18] K. Rohr, "Image registration based on thin plate splines and local estimates of anisotropic landmark localization uncertainties," in *Proc. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, pp. 1174–1183.
- [19] C. Davatzikos and R. N. Bryan, "Using a deformable surface model to obtain a shape representation of the cortex," *IEEE Trans. Med. Imag.*, vol. 15, pp. 785–795, Dec. 1996.
- [20] M. Vaillant and C. Davatzikos, "Hierarchical matching of cortical features for deformable brain image registration," in *Proc. Int. Conf. Information Processing in Medical Imaging*, June 1999, pp. 182–195.
- [21] H. Chui, L. Win, R. Schultz, J. Duncan, and A. Rangarajan, "A unified feature registration method for brain mapping," in *Proc. Int. Conf. Information Processing in Medical Imaging*, June 2001, pp. 300–314.
- [22] 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 *Proc. SPIE Conf. Geometrical Methods in Applied Imaging*, vol. 2573, July 1995, pp. 278–289.
- [23] J. P. Thirion, O. Monga, S. Benayoun, A. Gueziec, and N. Ayache, "Automatic registration of 3-D images using surface curvature," *Proc. SPIE Mathematical Methods in Medical Imaging*, vol. 1768, pp. 206–216, 1992.
- [24] M. Breijl and M. Sonka, "Object localization and border detection criteria design in edge-based image segmentation: Automated learning from examples," *IEEE Trans. Med. Imag.*, vol. 19, pp. 973–985, Oct. 2000.
- [25] R. Bajcsy, R. Lieberman, and M. Reivich, "A computerized system for the elastic matching of deformed radiographic images to idealized atlas images," *J. Comput. Assist. Tomogr.*, vol. 7, no. 4, pp. 618–625, 1983.

- [26] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. Image Processing*, vol. 5, Sept. 1996.
- [27] A. C. Evans, W. Dai, L. Collins, P. Neeling, and S. Marett, "Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis," *Proc. SPIE Image Processing*, vol. 1445, pp. 236–246, 1991.
- [28] P. T. Fox, M. A. Mintum, E. M. Reiman, and M. E. Reichle, "Enhanced detection of focal brain responses using inter-subject averaging and distribution analysis of subtracted PET images," *J. Cereb. Flow Metabol.*, vol. 8, pp. 642–653, 1988.
- [29] P. E. Roland, C. J. Graufelds, and J. Wahlin *et al.*, "Human brain atlas: For high-resolution functional and anatomical mapping," *Human Brain Mapping*, vol. 1, pp. 173–184, 1994.
- [30] J. P. Thirion, "Non-rigid matching using demons," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, June 1996, pp. 245–251.
- [31] D. Shen, E. H. Herskovits, and C. Davatzikos, "An adaptive-focus statistical shape model for segmentation and shape modeling of 3-D brain structures," *IEEE Trans. Med. Imag.*, vol. 20, pp. 257–270, Apr. 2001.
- [32] D. Shen, S. Moffat, S. M. Resnick, and C. Davatzikos, "Measuring size and shape of the hippocampus in MR images using a deformable shape model," *NeuroImage*, vol. 15, no. 2, pp. 422–434, Feb. 2001.
- [33] D. Shen and H. H. S. Ip, "Discriminative wavelet shape descriptors for invariant recognition of 2-D patterns," *Pattern Recogn.*, vol. 32, no. 2, pp. 151–165, Feb. 1999.
- [34] —, "Generalized affine invariant image normalization," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 19, pp. 431–440, May 1997.
- [35] D. Shen, H. H. S. Ip, K. K. T. Cheung, and E. K. Teoh, "Symmetry detection by generalized complex (GC) moments: A close-form solution," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 21, pp. 466–476, May 1999.
- [36] C. H. Teh and R. T. Chin, "On image analysis by the methods of moments," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, pp. 496–513, June 1988.
- [37] R. R. Bailey and M. Srinath, "Orthogonal moment features for use with parametric and nonparametric classifiers," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, pp. 389–399, Apr. 1996.
- [38] C. H. Lo and H. S. Don, "3-D moment forms: Their construction and application to object identification and positioning," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, pp. 1053–1064, Oct. 1989.
- [39] A. F. Goldszal, C. Davatzikos, D. L. Pham, M. X. H. Yan, R. N. Bryan, and S. M. Resnick, "An image-processing system for qualitative and quantitative volumetric analysis of brain images," *J. Comput. Assist. Tomogr.*, vol. 22, no. 5, pp. 827–837, Sept.–Oct. 1998.
- [40] R. Collobert and S. Bengio, "SVMtorch: Support vector machines for large-scale regression problems," *J. Machine Learning Res.*, vol. 1, pp. 143–160, 2001.
- [41] B. C. Vemuri, J. Liu, and J. L. Marroquin, "Robust multimodal image registration using local frequency representations," in *Proc. Int. Conf. Information Processing in Medical Imaging*, June 2001, pp. 176–182.
- [42] G. E. Christensen, "Consistent linear-elastic transformations for image matching," in *Proc. Int. Conf. Information Processing in Medical Imaging*, 1999, pp. 224–237.
- [43] H. J. Johnson and G. E. Christensen, "Landmark and intensity-based consistent thin-plate spline image registration," in *Proc. Int. Conf. Information Processing in Medical Imaging*, June 2001, pp. 329–343.
- [44] T. Wakahara and K. Odaka, "Adaptive normalization of handwritten characters using global/local affine transformation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 20, pp. 1332–1341, Dec. 1998.
- [45] J. G. Liu, F. H. Y. Chan, F. K. Lam, and H. F. Li, "A new approach to fast calculation of moments of 3-D gray level images," *Parallel Comput.*, vol. 26, no. 6, pp. 805–815, May 2000.
- [46] G. E. Christensen, S. C. Joshi, and M. I. Miller, "Individualizing anatomical atlases of the head," *Vis. Biomed. Comput.*, pp. 343–348, 1996.
- [47] Y. Wang and L. H. Staib, "Boundary finding with prior shape and smoothness models," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 22, pp. 738–743, 2000.
- [48] J. C. Gee, C. Barillot, L. L. Briquer, D. R. Haynor, and R. Bajcsy, "Matching structural images of the human brain using statistical and geometrical image features," in *Proc. SPIE Visualization in Biomedical Computing*, vol. 2359, 1994, pp. 191–204.
- [49] S. M. Resnick, A. F. Goldszal, C. Davatzikos, S. Golski, M. A. Kraut, E. J. Metter, R. N. Bryan, and A. B. Zonderman, "One-year age changes in MRI brain volumes in older adults," *Cereb. Cortex*, vol. 10, pp. 464–472, 2000.
- [50] D. L. Pham and J. L. Prince, "Adaptive fuzzy segmentation of magnetic resonance images," *IEEE Trans. Med. Imag.*, vol. 18, pp. 737–752, Sept. 1999.
- [51] J. Weese, G. P. Penney, P. Desmedt, T. M. Buzug, D. L. G. Hill, and D. J. Hawkes, "Voxel-based 2-D/3-D registration of fluoroscopy images and CT scans for image-guided surgery," *IEEE Trans. Inform. Technol. Biomed.*, vol. 1, pp. 284–293, Dec. 1997.
- [52] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: Application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, pp. 712–721, Aug. 1999.