

HAMMER: Hierarchical Attribute Matching Mechanism for Elastic Registration

Dinggang Shen¹ and Christos Davatzikos^{1,2,3}

¹Department of Radiology, ²Department of Computer Science, ³Center for Computer-Integrated Surgical Systems and Technology, Johns Hopkins University
Email: dgshen@cbmv.jhu.edu, hristos@rad.jhu.edu

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 remarkably high accuracy in superposition of images from different subjects, thus enabling very precise localization of morphological characteristics in population studies. There are two major novelties in the proposed algorithm. First, it uses an attribute vector, i.e. a set of geometric moment invariants that is defined on each voxel in an image, 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. This is a fundamental deviation of our method, referred to as 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 energy functions, which are constructed to have significantly fewer local minima. This is achieved by hierarchically selecting features that have distinct attribute vectors, thus drastically reducing ambiguity in finding correspondence. A number of experiments in this paper have demonstrated excellent performance.

1. 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 can be used for spatial normalization of functional images, for group analysis and statistical parametric mapping [1]. It is also used in computational anatomy as a means for measuring structure, by adapting an anatomical template to individual anatomies [2-6,14]. Finally, it is used as a means for image data mining in lesion-deficit studies, as well as in stereotaxic neurosurgery for mapping anatomical atlases to patient images [7,8]. Therefore, many image analysis methods 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 the

anatomical correspondences established manually, semi-automatically, or fully automatically on a number of distinct anatomical features. Such features are distinct landmark points [9], or a combination of curves and surfaces, such as sulci and gyri [4,10-13,28]. The second family of methods is based on volumetric transformations, which seek to maximize the similarity between an image and a template, and assume that the image and the template have been acquired with the same imaging protocol [1,2,14-17,27,29,30].

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 don't evaluate a matching criterion on every single voxel in an image, but rather 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.

In medical imaging, it is important to build deformable anatomical models that take into account the underlying anatomy, and not simply the similarity of image intensity. Toward that end, this was recognized in [18], which used 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 (i.e. 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 (i.e. two points that lie on the hippocampal boundary, but belong to different parts of the structure) [19]. In this paper, we build upon the ideas of the approach in [18], by defining attribute vectors in volumetric images, and using them for 3D warping. Specifically, we present a method that has the following two key novelties:

- 1) We use a sequence of successive approximations of the energy function being optimized, by lower dimensional 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 it increases the dimensionality of the energy function. The driving features are selected hierarchically according to how uniquely they can be identified and matched to 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, we use *geometric moment invariants (GMI's)* as a means for representing the geometric structure of the underlying anatomy. GMI's have been used very successfully in different computer vision application [20] 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 grey matter 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 similar image intensity. However, when examined at different scales, the precentral and postcentral gyri have different characteristics, hence anatomists have assigned them different names. A GMI at a particular scale is calculated by placing a spherical neighborhood around each voxel, and finding thirteen quantities that are invariant to rotation [21]. The radius of the neighborhood defines the scale of the GMI. We evaluate the GMI at different scales (i.e. for different neighborhood sizes), and concatenate the resulting GMI's into a long attribute vector. The idea of a GMI is demonstrated schematically in Fig. 1, where different local structures have different GMI's.

We present experiments with images from elderly subjects, which display several difficulties, including reduced tissue contrast, significant atrophy, and motion artifacts. We demonstrate results of exquisite accuracy (see Figs. 8, 9), which for the first time, to our knowledge, enable very accurate morphometric analyses at an extremely high resolution, which is practically limited only by the voxel dimensions.

2. Methods

2.1 General Formulation

In deformable registration, the key issue is to detect anatomical correspondences between two brain images. Some parts of the anatomy in the brain image can be identified relatively more reliably than others, such as

roots of sulci and crowns of gyri, since they have distinct attribute vectors. Therefore, it is reasonable that the deformation of the template be initially influenced by these structures. Most deformable models iteratively apply displacements to individual voxels. This can cause unrealistic deformations as individual points are pulled towards noisy edges. HAMMER applies a deformation to *subvolumes*, and evaluates the similarity of attribute vectors over the whole subvolume. This deformation strategy greatly helps our method avoid local minima. We now describe the methodology in detail.

Let's assume that $T(\mathbf{x})$ is the intensity of the voxel \mathbf{x} in the template, and $S(\mathbf{y})$ is the intensity of the voxel \mathbf{y} in the individual's image. Here, \mathbf{x} and \mathbf{y} are the 3D coordinates of the voxels in a volume \mathbf{V} . The displacement field $u(\mathbf{x})$ defines the mapping from the coordinate system of the template T to the subject S , while the transformation $h(\mathbf{x})=\mathbf{x}+u(\mathbf{x})$ defined the mapping that transforms the template T into the shape of the subject S . The inverse transformation of $h(\mathbf{x})$ is $h^{-1}(\mathbf{x})$. 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 energy function minimized by HAMMER is defined as follows:

$$E = \sum_{\mathbf{x} \in \mathbf{V}} \omega_T(\mathbf{x}) \left(\frac{\sum_{\mathbf{z} \in n(\mathbf{x})} \varepsilon(\mathbf{z}) (1 - m(\mathbf{a}_T(\mathbf{z}), \mathbf{a}_S(h(\mathbf{z}))))}{\sum_{\mathbf{z} \in n(\mathbf{x})} \varepsilon(\mathbf{z})} \right) + \sum_{\mathbf{y} \in \mathbf{V}} \omega_S(\mathbf{y}) \left(\frac{\sum_{\mathbf{z} \in n(\mathbf{y})} \varepsilon(\mathbf{z}) (1 - m(\mathbf{a}_T(h^{-1}(\mathbf{z})), \mathbf{a}_S(\mathbf{z})))}{\sum_{\mathbf{z} \in n(\mathbf{y})} \varepsilon(\mathbf{z})} \right) + \beta \sum_{\mathbf{x}} \|\nabla^2 u(\mathbf{x})\| \quad (1)$$

The first energy term sums over the differences of the attribute vectors of the template voxels at position \mathbf{x} , with their corresponding subject voxels at the position $h(\mathbf{x})$. $\mathbf{a}_T(\mathbf{z})$ is the attribute vector of the template voxel \mathbf{z} , while $\mathbf{a}_S(h(\mathbf{z}))$ is the attribute vector of the subject voxel at the position $h(\mathbf{z})$. \mathbf{z} is a voxel in the subvolume $n(\mathbf{x})$ of the template voxel \mathbf{x} . The function $m(\mathbf{a}_T(\mathbf{z}), \mathbf{a}_S(h(\mathbf{z})))$ measures the similarity of the attribute vectors of two voxels, the template voxel \mathbf{z} and the subject voxel at the position $h(\mathbf{z})$. The value of $m(\dots)$ ranges from 0 to 1. Therefore, $1-m(\dots)$ measures the difference of the two voxels. The integration of the differences of the attribute vectors $\mathbf{a}_T(\mathbf{z})$ and $\mathbf{a}_S(h(\mathbf{z}))$ over the whole subvolume $n(\mathbf{x})$ of the voxel \mathbf{x} , instead of individual voxel \mathbf{x} , enables our model to evaluate the similarity of whole parts of the anatomy. The parameter $\varepsilon(\mathbf{z})$ assigns higher weights to the boundary voxels \mathbf{z} , which yields a good matching of corresponding boundaries in the template and subject images. The term $\sum_{\mathbf{z} \in n(\mathbf{x})} \varepsilon(\mathbf{z})$ is used for normalization. The weighting parameter $\omega_T(\mathbf{x})$ determines the relative weight of each voxel \mathbf{x} in the template. If one voxel has distinct attribute

vectors and can be identified relatively more reliably than others, then this voxel is assigned higher weight and influences the initial deformation of the template relatively more than other points. By adaptively changing the weights in this energy function, our model is able to adaptively focus on the different parts of the anatomy during different deformation stages. The second energy term is defined similarly as the first, but it is defined on the subject domain. Since weightings $\omega_s(\mathbf{y})$ and $\omega_T(\mathbf{x})$ are determined by the image contents, 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* “template to subject”. This formulation makes our model robust to initialization and suboptimal solutions (see Fig. 4), and makes the transformation consistent [25]. Naturally, the introduction of this term increases the computational time of our model. In order to overcome this limitation, we selectively focus on a small part of major voxels in the subject image, by setting zero weightings $\omega_s(\mathbf{y})$ to all other unreliable voxels. Although this does not warranty consistency of the full 3D transformation [25], it warranties consistency of the driving forces, something that 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.




Fig. 1 Demonstration of the GMI's in discriminating local structures. The GMI's of the voxel indicated by a black circle in (a) is compared with the GMI's of all other points. The resulting similarity of the GMI's is shown in (b), where the white circle corresponds to the black circle in (a). High similarity is white.

2.2 Attribute Vectors

As we described in the introduction, the attribute vectors are a critical part of our formulation. Consider a voxel with 3D coordinates \mathbf{x} in a volumetric image. Then the attribute vector, $\mathbf{a}(\mathbf{x})$, at \mathbf{x} is defined as

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

where the individual components of the attribute vector are the following. The first component $\mathbf{a}_1(\mathbf{x})$ is the edge type of the voxel \mathbf{x} . There exist several edge types, i.e. non-edge and six combinations of edges among grey matter (GM), white matter (GM), and CSF. This attribute can be obtained directly after performing intensity segmentation

on the brain images [22], or it can be determined by using learning algorithm such as a support vector machine [23] from the brain samples. The second component $\mathbf{a}_2(\mathbf{x})$ is the intensity of the voxel \mathbf{x} . The last part of the attribute vector $\mathbf{a}_3(\mathbf{x})$ comprises the GMI's at different scales. For each scale, there are thirteen rotation invariants that are calculated from the zero-order, second-order, third-order 3D regular moments [21]. In order to calculate the rotation invariants for the voxel \mathbf{x} , let's assume that the origin of the coordinate system has been shifted to the voxel \mathbf{x} , and under this coordinate system the intensity function of the volumetric brain image is $f(x_1, x_2, x_3)$, where x_1, x_2, x_3 are coordinates. Then, the 3D regular moments of order $(p+q+r)$ of the image $f(x_1, x_2, x_3)$ can be defined by

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

where R is the size of the neighborhood around the origin (i.e. voxel \mathbf{x}). Two rotation invariants are listed next. Other eleven rotation invariants can be found in [21].

$$I_1 = M_{0,0,0}, \quad I_2 = M_{2,0,0} + M_{0,2,0} + M_{0,0,2}$$

If we can use intensity segmentation method to label brain tissue into GM, WM, and CSF, then we can calculate thirteen GMI's for each tissue, with the first GMI corresponding to the volume of this tissue. In this way, we have 13x3 GMI's for each voxel at each scale.

GMI's are a convenient way of characterizing the shapes of objects in 2D or 3D images, since objects of different shapes tend to have different GMI's. Fig. 1 demonstrates the concept of GMI's that can be used to discriminate local structures. In this figure, the GMI's of the point indicated by a black circle in Fig. 1a are compared with the GMI's of all other points in the image. Fig. 1b shows the attribute vector similarity of the selected point and the other points, ranging from 0 to 1. The size of this image is 256x256, while the GMI's are calculated in 81x81 neighborhood. It should be indicated that, besides using GMI's to extract the geometric properties of objects, other techniques such as Gabor filters, Wavelets, and low-frequency representations [24], can also be applied, under the condition that the extracted attributes are rotation invariant.

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. This is because boundaries are very important features in describing the brain structures. Mathematically, the similarity of the two voxels \mathbf{x} and \mathbf{y} is defined as

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

where $c(\cdot)$ is the similarity of the second and third parts of attribute vectors. When a set of training samples is available, the similarity $c(\cdot)$ can be learned by a learning technique such as support vector machines [23]. Without the set of training samples, we can simply normalize each

element in \mathbf{a}_2 and \mathbf{a}_3 to a range from 0 and 1, and then define $c(.,.)$ as

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

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

2.3. Image Deformation Mechanism: hierarchical approximation of the energy function

In order to optimize the criterion in (1), we have developed a deformation mechanism that is robust to suboptimal solutions. This deformation mechanism is based on the approximation of the energy function in (1), by an energy function, with the significantly fewer local minima, which is defined on the *focus points* in the template and subject, i.e. points that drive the volumetric deformation procedure at a particular stage. Initially, only few distinctly identifiable features are used to drive the deformation, as described below.

Focus points. 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. Good examples are the roots of sulci and the crowns of the gyri, which can be identified 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.

Our deformation mechanism addresses exactly the issues above. In particular, certain parts of the anatomy are used first to drive the volumetric deformation. These parts include sulcal roots and gyral crowns, as well as regions in which edges are relatively stronger and more isolated, i.e. ventricular boundaries and outside cortical boundaries. It is particularly important in our studies to use the ventricular boundaries to drive the volumetric deformation in the very beginning stages, since this way we account for large image dissimilarities caused by the extreme atrophy of many of our elderly subjects.

Specifically, the rules for the hierarchical procedure of our deformation mechanism are the following:

1) Roots of sulci and crowns of gyri are typically identified very robustly, based on their distinctive attribute vectors. This is demonstrated in Fig. 2(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 ones.

2) Regions that display relatively higher 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 higher.

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 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 [18].

Fig. 2 gives an example of the focus voxels during different stages of the deformation procedure that deforms the template in (a1-a4) to the subject in (b). The subject image in 2b is the same as the cross-section image in Fig. 6(Li), which is automatically affine-transformed from the initial input subject (in Fig. 6(I)) by HAMMER. The black points in 2(a1) are the voxels that our model automatically determined for the initial deformation stages. Most of them are the voxels in the sulcal roots and gyral crowns. (Only points on the grey/white matter interface are shown, so 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 subject, other boundary voxels are added to the set of the focus voxels and immediately join the deformation procedure. For example, the white voxels are added in 2(a2), then other black and white voxels are added in 2(a3) and 2(a4). The black voxels in 2b are voxels in the subject that have distinctive attribute vectors. Only points on the white matter/grey matter boundary are shown, for clarity. These points tend to fall on sulci or gyri, which are easier to distinguish based on their attribute vectors.




Fig. 2 Schematic demonstration of the hierarchical deformation mechanism in HAMMER. See the text for details.

Displacement of subvolumes. The vast majority of deformable registration methods rely on iterative methods that effectively move one voxel at a time, trying to improve some similarity criterion. However, the huge dimensionality of systems arising in 3D warping methods of medical images, in conjunction with the highly non-convex 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.

In our approach, we have chosen a different deformation mechanism, which has proven to be very robust to local minima. In particular, for each template voxel that is considered at a particular stage of the deformation, a subvolume, i.e. a small volume around that voxel, is displaced at several tentative nearby locations, based on which nearby voxels have similar attribute vectors with the voxel under consideration. *Importantly, the optimal deformation of the subvolume is not determined based on the similarity of one voxel, but it is determined by integrating the similarity of the attribute vectors within the whole subvolume.* For example, 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. Otherwise, this voxel will produce no deformation on its subvolume. This is shown schematically in Fig. 3. Figs. 3b and 3a respectively display the selected part of the template and the subject. A regular grid \mathbf{x} in the template (b) has been deformed slightly to an irregular grid $h(\mathbf{x})$ in the subject (a). The white “disc” in Fig 3 denotes a template voxel under consideration. The search domain of this voxel is shown as the big dotted circles in c1 and d1. Two voxels with the similar attribute vectors are found in its neighborhood as shown as white circles in c1 and d1. Then the subvolume of this voxel is tentatively deformed to these two positions, as shown in c2 and d2. The corresponding deformed template grids $h(\mathbf{x})$ are given in c1 and d1. By integrating the total similarity of the attribute vectors in the current deformed subvolume with those in the subject, this subvolume is finally deformed to an optimal position such as c2. White crosses in c2 and d2 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 could be 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 can be subsampled according to the size of this subvolume.




Fig. 3 Demonstration of the deformation mechanism in HAMMER. See the text for detailed meanings of the pictures.

The mathematical form of the deformation of the subvolume is described next. When deforming the subvolume of the template voxel \mathbf{x} , we should produce no discontinuities at the boundary of the subvolume. In order to reach this objective, we propagate the deformation on the voxel \mathbf{x} to its neighboring voxels \mathbf{z} according to a Gaussian kernel. Currently, the position of the template voxel \mathbf{x} is $h(\mathbf{x})$, and the position of the voxel \mathbf{z} that is included in the subvolume of the voxel \mathbf{x} is $h(\mathbf{z})$. Assume that the position of template voxel \mathbf{x} is moved from $h(\mathbf{x})$ to a tentative position $h(\mathbf{x}) + \Delta$, during the greedy research. Then, the new position of the voxel \mathbf{z} is defined as

$$h(\mathbf{z}) + \Delta \cdot \exp(-\|\mathbf{x} - \mathbf{z}\|^2 / 2\sigma^2),$$

where σ is a parameter that makes $\exp(-\|\mathbf{x} - \mathbf{z}\|^2 / 2\sigma^2)$ close to zero for the voxels \mathbf{z} in the boundary of the subvolume around the template voxel \mathbf{x} . This definition leaves the boundary voxels unchanged and hence maintains the smoothness of the displacement field.

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 2 years [25]. The similar concept of consistent transformations has been previously introduced in the computer vision community by [26], for normalization of hand-written characters. In this paper, we used a consistency term in eq. (1), which constrains the inverse transformation $h^{-1}(\mathbf{x})$ and pertains to the deformation of “subject to template”. This formulation makes our model very robust to the initialization and suboptimal solutions. Fig. 4 vividly demonstrates the importance of using consistency term, with a good warping result using the consistency term in Fig. 4a and an incorrect warping result, without using the consistency term, in Fig. 4b. The black contours in Fig. 4 are the WM boundaries of the template brain image in the same slice.




Fig. 4 Demonstration on the importance of using the consistency term in our energy function. (a) Using the consistency term, (b) without using the consistency term.




Fig. 5 Schematic description of our idea in converting the constraint on the inverse transformation $h^{-1}(\mathbf{x})$ as an additional constraint on the transformation $h(\mathbf{x})$. See text for details.

In the numerical implementation, even if we know the transformation $h(\mathbf{x})$, it is still computationally burdensome to determine the inverse transformation $h^{-1}(\mathbf{x})$. We bypass the problem of finding the inverse transformation by converting the constraint on the inverse transformation $h^{-1}(\mathbf{x})$ as an additional constraint on the transformation $h(\mathbf{x})$. Moreover, in order to speed up our algorithm, we only focus on the major subitems in the second energy term of eq. (1), such as maintaining only subitems with big weightings $\omega_s(\mathbf{y})$ that are corresponding to the distinctive boundary voxels in the subject, and removing all other subitems.

Assume that, in the current iteration, the set of the template focus voxels \mathbf{x}_i is $\{\mathbf{x}_i | 1 \leq i \leq N_T\}$, such as a number of small black circles in Fig 5a. Here, N_T is the number of the template focus voxels. Similarly, let $\{\mathbf{y}_i | 1 \leq i \leq N_S\}$ be the set of the N_S major subject image's focus voxels. This latter set of focus voxels can be selected adaptively, as described earlier for the template. However, for saving computation time, in the current study we fix this set in the whole deformation procedure. For example, Fig 2b shows the subject's focus voxels in a slice. With the transformation function $h(\mathbf{x})$, we know the set of the template voxels $\{\mathbf{x}_i | 1 \leq i \leq N_T\}$ has been displaced into $\{h(\mathbf{x}_i) | 1 \leq i \leq N_T\}$, such as small black circles in Fig 5b. Then for each subject voxel \mathbf{y}_i under consideration, such as a grey point, in its neighborhood we find a displaced template voxel $h(\mathbf{x}_j)$, where $h(\mathbf{x}_j) \in \{h(\mathbf{x}_i) | 1 \leq i \leq N_T\}$, with the most similar attribute vector. If the similarity between the subject voxel \mathbf{y}_i and the template voxel \mathbf{x}_j is high, then a force is applied on the displaced template voxel $h(\mathbf{x}_j)$ at the

direction from $h(\mathbf{x}_j)$ to \mathbf{y}_i , as shown as a grey arrow in Fig 5b. Finally, the displaced template voxel $h(\mathbf{x}_j)$, including its subvolume, will be deformed according to the definition of the displacement of subvolume, described earlier.

Multi-grid formulation. In order to speed up HAMMER, we have implemented it in a multi-grid framework. Currently, we use 3 different scales, corresponding to the original image resolution (High), a subsampled version by a factor of 2 (Middle), and a subsampled version by a factor of 4 (Low). For each resolution, we first calculate the attribute vector for each voxel in the template and subject. Then we use HAMMER to register them. The transformation function $h(\mathbf{x})$ in one low resolution are upsampled and linearly interpolated to the next resolution. The upsampled transformation function is then used as the initialization for the studied resolution. An example of warping two brains using this multi-grid formulation is demonstrated in Fig. 6 in Section 3.

3. Results

Two experiments are given here to demonstrate the performance of HAMMER. As to the speed of HAMMER, it currently takes about 4~5 hours on an SGI OCTANE workstation for a full 3D warping of MR images, without any code optimization.

Experiment on warping one subject to the template. Fig 6 shows the result obtained by warping the subject with large ventricles to the template. The black contours in Fig 6 are the WM boundaries of the template brain image at the same slice, which are used as a reference for comparison. The small images in the upper left corner were taken from the low resolution warping procedure. In the beginning, the subject image was automatically transformed to the space of the template by using the global affine transformation. A representative slice image of this affine transformed volume is displayed in 6(Li). The effect of this affine transformation can be observed by comparing three images in 6(I), (Li), and (Lt). Notice that an image in 6(I) is the original subject image, and an image in 6(Lt) is the template image. All of them are at the same slice. The final warping result of the low resolution procedure is under the image 6(Lt). We can observe that the subject has been warped a lot to the template. By upsampling the displacement fields from the low resolution to the middle resolution, we can obtain the initial warping image in the middle resolution, as shown in 6(Mi). This warping result was improved in the end of the middle resolution, by comparing the template in 6(Mt) with the warping result under this template. Similarly, the displacement fields were upsampled and interpolated to the high resolution, and we therefore obtained the initial warping result in this high resolution in 6(Hi). The final warping result of the subject is displayed in the right

corner. Comparing this final warping result with the template in 6(Ht), we can easily observe the exquisite accuracy of HAMMER in warping brain images. A movie on the procedure of warping this subject can be found in the web site:

http://beast.cbm.v.jhu.edu/~dgshen/HAMMER/Whole_withCaption.gif.

In Fig 7, we provide a picture of the displacement fields of this warping. We can observe the relative smoothness of the associated displacement fields, even the initial difference between the ventricles of the subject and the template is large.




Fig. 6 A procedure of warping a subject brain to the template. See text for details. Bottom right shows warped subject overlaid on template's outline




Fig. 7. Displacement fields on a typical slice of the template.

Experiment on constructing average brain image. The quality of the average brain image can be used to validate the accuracy of the warping algorithm in performing registration. We selected 18 subjects, with a wide variability of ventricular shape and size, and, of course representative of variability in other structures. We applied HAMMER to warp these 18 subjects using the same fixed set of parameters. Fig. 8 shows an average brain image in

three different views. For comparison, the same slice images of the template are also displayed. Notably, the cortical regions, ventricles, and WM/GM boundaries are very clear. Also, the caudate nucleus and the lenticular nucleus are also very clear. That indicates the high registration accuracy obtained by HAMMER.

We also generated a 3D rendering of the average image, shown in Fig 9 along with the template brain. An excellent definition of the cortical anatomy is evident in this average.




Fig. 8 Average brain image reconstructed from 18 very different subjects, along with the template slices images. The crispness of the average reflects an exquisite registration accuracy.




Fig. 9 The 3D renderings of the 18-brain average brain and the template brain. The black crosses are used for comparing the correspondences in the average and template brains.

4. Conclusion

We presented a new approach to deformable registration of brain images, which resulted in extremely high accuracy in superposition of images from different subjects, thus enabling very precise localization of morphological characteristics in population studies. The proposed approach has two major novelties. First, it utilizes attribute vectors to reflect geometric characteristics of the underlying anatomical structures. The attribute vectors are essential in establishing anatomical correspondences in the deformation procedure. Second, it uses a hierarchical deformation mechanism. In particular, parts of the anatomy that have distinct shape properties are used as “anchors” during initial deformation stages, while the remaining parts simply follow the deformation. As a template deforms toward its target configuration, other parts of the anatomy gradually affect the deformation, as those get close to their respective targets. Moreover, the deformation mechanism deforms subvolumes at a time, rather than individual voxels, and evaluates the similarity of attribute vectors over the entire subvolume.

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 exquisite crispness of the resulting averages after warping (see Figs. 8,9), along with the relative smoothness of the associated displacement fields (see Fig. 7), are very promising, in that they will allow us to examine morphology at a very detailed level, compared to the image resolution.

References:

1. K.J. Friston, A.P. Holmes, K.J. Worsley, J.P. Poline, C.D. Frith, 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 and M. Reivich and R. Bajcsy. “Elastically Deforming 3D Atlas to Match Anatomical Brain Images”, *J. Comp. Assist. Tomogr.*, volume 17, p. 225-236, 1993.
3. C. Davatzikos. “Mapping of image data to stereotaxic spaces: applications to brain mapping”. *Human Brain Mapping*, 6:334-338, 1998.
4. P. Thompson and A.W. Toga. “A surface-based technique for warping three-dimensional images of the brain”. *IEEE Trans. on Med. Imaging*, volume 15, p.402-417, 1996.
5. D.L. Collins, P. Neelin, T.M. Peters and A.C. Evans. “Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space”. *J. of Comp. Ass. Tom.*, Vol. 18, p. 192-205, 1994.
6. F.L. Bookstein. “Thin-plate splines and the atlas problem for biomedical images”. *IPMI’91*, p. 326-342, 1991.
7. S.K. Kyriacou, C.A. Davatzikos, S.J. Zinreich, and R.N. Bryan. “Nonlinear Elastic Registration of Brain Images with Tumor Pathology Using a Biomechanical Model”. *IEEE Transactions on Medical Imaging*, Vol. 18(7), pp. 580-592.
8. T. Peters, B. Davey, P. Munger, R. Comeau, A. Evans and A. Olivier. “Three-dimensional multimodal image guidance for neurosurgery”. *IEEE Trans. on Med. Imaging*, p.121-128, 1996.
9. K. Rohr. “Image registration based on thin plate splines and local estimates of anisotropic landmark localization uncertainties”. *MICCAI’98*, volume 1496, p.1174-1183, 1999.
10. M. Vaillant and C. Davatzikos. “Hierarchical Matching of Cortical Features for Deformable Brain Image Registration”. *IPMI’99*, Volume 1613, p.182-195, June 1999.
11. H. Chui, L. Win, R. Schultz, J. Duncan, and A. Rangarajan. “A Unified Feature Registration Method for Brain Mapping”. *IPMI’01*, p.300-314, Davis, CA, USA, June 18-22, 2001.
12. 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”. *Proc. of the SPIE Conf. on Geom. Methods in Applied Imaging*, vol. 2573, p.278-289, July 1995.
13. M. Breijl and M. Sonka. “Object localization and border detection criteria design in edge-based image segmentation: automated learning from examples”. *IEEE Trans. on Med. Imaging*, vol. 19, p.973-985, 2000.
14. G.E. Christensen, R.D. Rabbitt and M.I. Miller. “Deformable Templates Using Large Deformation Kinematics”. *IEEE Trans. on Image Processing*, volume 5, number 9, September 1996.
15. 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”. *SPIE Proc., Image Processing*, volume 1445, p.236-246, 1991.
16. 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. Cerebral Flow and Metabolism*, volume 8, p.642-653, 1988.
17. P.E. Roland, C.J. Graufelds, J. Wahlin, et al. “Human Brain Atlas: for High-Resolution Functional and Anatomical Mapping”. *Human Brain Mapping*, volume 1, p.173-184, 1994.
18. D. Shen, E.H. Herskovits, C. Davatzikos. “An adaptive-focus statistical shape model for segmentation and shape modeling of 3D brain structures”. *IEEE Trans. on Med. Imag.*, 20(4):257-270, 2001.
19. 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*, submitted.
20. 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. on PAMI*, 21(5):466-476, May 1999.
21. C.H. Lo, H.S. Don. “3-D Moment Forms: Their Construction and Application to Object Identification and Positioning”. *IEEE Trans. on PAMI*, 11(10): 1053-1064, Oct. 1989.
22. A.F. Goldszal, C. Davatzikos, D.L. Pham, M.X.H. Yan, R.N. Bryan, S.M. Resnick. “An image-processing system for qualitative and quantitative volumetric analysis of brain images”. *Journal of Computer Assisted Tomography*, 22 (5): 827-837, Sept.-Oct. 1998.
23. R. Collobert and S. Bengio. “SVM-Torch: Support Vector Machines for Large-Scale Regression Problems”. *Journal of Machine Learning Research*, vol. 1, p.143-160, 2001.
24. B. C. Vemuri, J. Liu, and J.L. Marroquin. “Robust Multimodal Image Registration Using Local Frequency Representations”. *IPMI’01*, p.176-182, Davis, CA, USA, June 18-22, 2001.
25. H.J. Johnson and G.E. Christensen. “Landmark and Intensity-Based, Consistent Thin-Plate Spline Image Registration”. *IPMI’01*, p.329-343, Davis, CA, USA, June 18-22, 2001.
26. T. Wakahara, K. Odaka. “Adaptive normalization of handwritten characters using global/local affine transformation”. *IEEE Trans. on PAMI*, 20(12): 1332-1341, Dec. 1998.
27. Y. Wang and L.H. Staib. “Boundary finding with prior shape and smoothness models”. *IEEE Trans. on PAMI*, p. 738-743, 2000.
28. J.P. Thirion, O. Monga, S. Benayoun, A. Gueziec and N. Ayache. “Automatic registration of 3D images using surface curvature”. *SPIE Proc., Math. Methods in Med. Imaging*, vol. 1768, p.206-216, 1992.
29. J.P. Thirion. “Non-rigid matching using demons”. *Proc. IEEE Conf. Comp. Vis. and Patt. Recog.*, 1996.
30. 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”. *Proc. SPIE Visualization in Biomedical Computing 1994*, vol. 2359, pp.191-204, 1994.