

Measuring temporal morphological changes robustly in brain MR images via 4-dimensional template warping

Dinggang Shen* and Christos Davatzikos

Section of Biomedical Image Analysis, Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104-2644, USA

Received 5 August 2003; revised 6 December 2003; accepted 9 December 2003

Robustly measuring subtle longitudinal brain changes is a major challenge in computational neuroanatomy. This paper describes a method for measuring temporal morphological brain changes, by means of a 4-dimensional image warping mechanism. Longitudinal stability is achieved by considering all temporal MR images of an individual simultaneously, rather than by individually warping a template to an individual, or by warping the images of one time-point to those of another time-point. Following earlier work in 3D, a local morphological signature is attached to each voxel of a sequence of images, and it includes a set of image attributes reflecting morphological characteristics of the spatiotemporal structure around the respective voxel at different scales. This attribute vector forms the basis for searching in the 4-dimensional space for a counterpart that has similar morphological signature, thereby leading to automated detection of anatomical correspondence. Ambiguities in this process are reduced by constructing attribute vectors that are highly distinctive of respective voxels, and by using a hierarchical matching procedure in which reliable and easily distinguishable voxels are used to guide the 4D deformation process. The resultant deformations are smooth both in the spatial and temporal dimensions, and are shown to significantly improve warping accuracy over a series of independent 3D warpings, in longitudinal measurements.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Atlas warping; Longitudinal change; Temporal change; Brain atrophy; Brain morphology; Deformable registration; Attribute vectors; Hierarchical deformation; 4D registration

Introduction

Measuring longitudinal morphological brain changes requires highly accurate segmentation and volumetric measurement of brain structures, either manually or by labeled atlas warping, from images acquired at different time-points (Chung et al., 2001). Manual segmentation requires extensive human interaction and

considerable training. Most importantly, both intra-rater reproducibility and inter-rater agreement are difficult to achieve in a longitudinal study, particularly when small changes are to be measured. Accordingly, many automatic image segmentation and parcellation methods, often based on atlas matching and registration, have been developed (Bajcsy et al., 1983; Bookstein, 1989; Breijl and Sonka, 2000; Christensen et al., 1996; Chui et al., 2001; Chung et al., 2001; Davatzikos and Bryan, 1996; Evans et al., 1991; Fox et al., 1988; Friston et al., 1995; Gee et al., 1994; Joshi et al., 1995; Miller et al., 1993; Rohr, 1999; Roland et al., 1994; Rueckert et al., 1999; Thirion, 1996; Thirion et al., 1992; Thompson and Toga, 1996; Vaillant and Davatzikos, 1999; Weese et al., 1997). Automatic registration methods generally fall into two categories. The first category includes methods based on feature matching, in which the spatial transformations are calculated from a number of anatomical correspondences established manually, semi-automatically, or fully automatically for a number of distinct anatomical features (Bookstein, 1989; Breijl and Sonka, 2000; Chui et al., 2001; Davatzikos and Bryan, 1996; Joshi et al., 1995; Miller et al., 1993; Rohr, 1999; Thirion et al., 1992; Thompson and Toga, 1996; Vaillant and Davatzikos, 1999). Such features are distinct landmark points (Bookstein, 1989; Rohr, 1999), or a combination of curves and surfaces, such as sulci and gyri (Breijl and Sonka, 2000; Chui et al., 2001; Davatzikos and Bryan, 1996; Joshi et al., 1995; Thirion et al., 1992; Thompson and Toga, 1996; Vaillant and Davatzikos, 1999). The second category includes methods based on volumetric transformations, which seek to maximize the similarity between a subject and a template, and generally assume that the subject and the template have been acquired by the same imaging protocol (Bajcsy et al., 1983; Christensen et al., 1996; Evans et al., 1991; Fox et al., 1988; Friston et al., 1995; Gee et al., 1994; Roland et al., 1994; Rueckert et al., 1999; Thirion, 1996; Weese et al., 1997). We previously proposed a fully automated atlas matching approach, referred to as Hierarchical Attribute Matching Mechanism for Elastic Registration (HAMMER) (Shen and Davatzikos, 2002, 2003), which uses a hierarchical morphological attribute-based deformation strategy. That approach attempted to combine advantages from both of the categories described above, by treating each voxel as a feature characterized by its own morphological signature, namely its *attribute vector*.

Regardless of their relative merits and limitations, all of the above-mentioned atlas warping techniques were mainly designed

* Corresponding author. Department of Radiology, University of Pennsylvania, 3600 Market Street, Suite 380, Philadelphia, PA 19104-2644. Fax: +1-215-614-0266.

E-mail addresses: dgshen@rad.upenn.edu (D. Shen), christos@rad.upenn.edu (C. Davatzikos).

Available online on ScienceDirect (www.sciencedirect.com.)

for 3D images. Consequently, applying these 3D methods independently for each time-point in a longitudinal study typically leads to noisy longitudinal measurements, particularly for small structures such as the hippocampus, due to inconsistencies in atlas matching among different time-points. It is usually assumed that the longitudinal measurements are smooth along the temporal dimension, as long as the longitudinal images of the subject are collected with adequate temporal resolution. It is particularly true for one of our studies, namely the Baltimore Longitudinal Study of Aging (BLSA), where consecutive longitudinal MR images of older adults were acquired within a relatively short interval of about a year. Although one can always obtain a smooth estimation of longitudinal changes by smoothing the measurements along the temporal dimension, the smoothed measurements will, in general, deviate significantly from the actual image data, unless smoothing is performed concurrently with atlas warping and by taking into account the image features. Freeborough and Fox (1998) imposed temporal smoothing by warping the images of one time-point to those of another time-point. However, that approach, too, does not avoid jitter noise between any two differential measurements, since spatial normalization is based on independent warpings between pairs of images. Most importantly, when these longitudinal correspondences are warped to the space of a template for automatic labeling or group analyses, they are subject to registration errors.

In this paper, we present an approach that overcomes this limitation, by simultaneously establishing the longitudinal correspondences in the individual and the inter-subject correspondences between the template and the individual. This is different from the 3D warping methods, which aim at establishing only the inter-subject correspondences between the template and the individual in a single time-point. Specifically, we propose a fully automatic 4-dimensional atlas matching method that constrains the smoothness in both spatial and temporal domains during the hierarchical atlas matching procedure, thereby producing smooth and accurate estimations of longitudinal changes. Most importantly, morphological features and matches guiding this deformation process are determined via 4D image analysis, which significantly reduces noise and improves robustness in detecting anatomical correspondence. Put simply, image features that are consistently recognized in all time-points guide the warping procedure, whereas spurious features, such as noisy edges, appear inconsistently at different time-points and are eliminated.

Methods

Overview

The proposed approach, referred to as 4D-HAMMER, involves the following two steps:

(1) Rigid alignment of 3D images of a given subject acquired at different time-points, to produce a 4D image. We employ 3D-HAMMER to establish the correspondences between neighboring 3D images, and then align one image (time t) to its previous-time image ($t - 1$) by a rigid transformation calculated from the established correspondences (Umeyama, 1991). This step simply removes variations in patient positioning in the scanner.

(2) Hierarchical deformation of the 4D atlas to the 4D subject images, via a hierarchical attribute-based matching method. Initially, the deformation of the atlas is influenced primarily by voxels with distinctive attribute vectors, thereby minimizing the chances of poor matches and also reducing computational burden. As the deformation proceeds, voxels with less distinctive attribute vectors gradually gain influence over the deformation. The resultant deformation fields are smooth in both spatial and temporal domains, via appropriate constraints defined in both domains.

Notation

The 4D images studied in this paper have three spatial dimensions (\mathbf{x}) and one temporal dimension (t), where \mathbf{x} and t are the spatial and the temporal coordinates, respectively. Let $T(\mathbf{x}, t)$ denote the intensity of the 4D template image at a point (\mathbf{x}, t) ; therefore, $T_t(\mathbf{x})$ is a usual 3D template image at time t . Similarly, $S(\mathbf{x}, t)$ denotes the intensity of the 4D subject image, and $S_t(\mathbf{x})$ is a usual 3D subject image at time t . The spatial domains for the template and the subject are not necessarily the same; hence, V_T is used for the template's spatial domain and V_S for the subject's spatial domain. However, the temporal domains are assumed to be same, that is, totally N volumetric images constitute the template and the subject, where $t \in [1, N]$. Since the 3D images are acquired at known time-points and temporal correspondence is thus known, there is no displacement along the temporal dimension during the deformable registration procedure. That is, the template image $T_t(\mathbf{x})$ is only permitted to warp toward the subject image of the same time t , $S_t(\mathbf{x})$. Accordingly, the displacement field $u(\mathbf{x}, t)$ defines the relative *spatial* deformation of the template image $T(\mathbf{x}, t)$ to the subject image $S(\mathbf{x} + u(\mathbf{x}, t), t)$. Let $h(\mathbf{x}, t)$ be the transformation corresponding to the displacement field $u(\mathbf{x}, t)$: $h(\mathbf{x}, t) = (\mathbf{x} + u(\mathbf{x}, t), t)$. The inverse transformation is $h^{-1}(\mathbf{x}, t)$, which will be referred as a backward transformation in the next.

Attribute vector

As it was described in the introduction, each voxel carries its own morphological signature, that is, an attribute vector $\mathbf{a}(\mathbf{x}, t)$, which is designed to be as distinctive as possible of its respective voxel, to facilitate automated matching. Each attribute vector includes edge types, image intensity, and Geometric Moment Invariants (GMIs) (Lo and Don, 1989), all of which are computed from the 3D spatial images (Shen and Davatzikos, 2002). 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 (Lo and Don, 1989). 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 reason for developing an attribute vector is that, if rich enough, the attribute vector can distinguish between different parts of the image. The detailed definitions for the attribute vectors and the similarity definitions of those vectors can be found in our previous paper (Shen and Davatzikos, 2002). Even richer geometric features, such as those based on Wavelet coefficients, can also be used as image attributes, as we described in Xue et al. (submitted for publication).

Mathematical optimization

The 4D-HAMMER algorithm hierarchically minimizes the cost function, E , of attribute vector similarity, which is given below:

$$E = E_F + E_B + E_S, \quad (1)$$

where

$$E_F = \sum_{t=1}^N \sum_{\mathbf{x} \in V_T} \omega_T(\mathbf{x}, t) \times \left(\frac{\sum_{(\mathbf{z}, \tau) \in n(\mathbf{x}, t)} \varepsilon_T(\mathbf{z}, \tau) (1 - m(\mathbf{a}_T(\mathbf{z}, \tau), \mathbf{a}_S(h(\mathbf{z}, \tau))))}{\sum_{(\mathbf{z}, \tau) \in n(\mathbf{x}, t)} \varepsilon_T(\mathbf{z}, \tau)} \right), \quad (2)$$

$$E_B = \sum_{t=1}^N \sum_{\mathbf{x} \in V_S} \omega_S(\mathbf{x}, t) \times \left(\frac{\sum_{(\mathbf{z}, \tau) \in n(\mathbf{x}, t)} \varepsilon_S(\mathbf{z}, \tau) (1 - m(\mathbf{a}_T(h^{-1}(\mathbf{z}, \tau)), \mathbf{a}_S(\mathbf{z}, \tau)))}{\sum_{(\mathbf{z}, \tau) \in n(\mathbf{x}, t)} \varepsilon_S(\mathbf{z}, \tau)} \right), \quad (3)$$

$$E_S = \alpha E_S^{\text{Spatial}} + \beta E_S^{\text{Temporal}}, \quad (4)$$

$$E_S^{\text{Spatial}} = \sum_{t=1}^N \sum_{\mathbf{x} \in V_T} \left\| u(\mathbf{x}, t) - \frac{\sum_{\mathbf{z} \in n(\mathbf{x})} g_1(\|\mathbf{z} - \mathbf{x}\|) u(\mathbf{z}, t)}{\sum_{\mathbf{z} \in n(\mathbf{x})} g_1(\|\mathbf{z} - \mathbf{x}\|)} \right\|^2, \quad (5)$$

$$E_S^{\text{Temporal}} = \sum_{\mathbf{x} \in V_T} \sum_{t=1}^N \left\| u(\mathbf{x}, t) - \frac{\sum_{\tau \in n(t)} g_2(\tau - t) u(\mathbf{x}, \tau)}{\sum_{\tau \in n(t)} g_2(\tau - t)} \right\|^2. \quad (6)$$

The first term E_F measures the similarity of attribute vectors in the vicinity of a point (\mathbf{x}, t) under consideration, and it is therefore guiding the identification of anatomical correspondence, to the extent that anatomical correspondence is reflected by similarity of the voxel-wise signatures. E_F is based on the forward transformation $h(\cdot)$ from the template to the subject. In particular, for each template point (\mathbf{x}, t) , we integrate the attribute vector differences in its 4D (spatio-temporal) neighborhood, denoted by $n(\mathbf{x}, t)$. (\mathbf{z}, τ) is a neighboring point. $\mathbf{a}_T(\mathbf{z}, \tau)$ is the attribute vector of the template point (\mathbf{z}, τ) , and $\mathbf{a}_S(h(\mathbf{z}, \tau))$ is the attribute vector of the corresponding subject point $h(\mathbf{z}, \tau)$. $m(\cdot, \cdot)$ is the similarity of two attribute vectors, and ranges from 0 to 1. Thus, $(1 - m(\cdot, \cdot))$ is the difference of two attribute vectors. $\varepsilon_T(\mathbf{z}, \tau)$ is the weighting parameter for the points in the neighbor-

hood; large weights are used for the salient points, such as boundary points. $\omega_T(\mathbf{x}, t)$ denotes the importance of the template point (\mathbf{x}, t) in image matching. Large weights are also assigned to the points with the distinctive attribute vectors. This assignment allows our algorithm to focus initially only on the distinctive points. By hierarchically assigning weights to the template points, our approach is able to hierarchically deform the images. Moreover, the size of neighborhood $n(\mathbf{x}, t)$ is large initially and decreases gradually with the progress of the deformation, which increases the robustness and accuracy of our algorithm. The second term E_B is similar to the first one, but it is defined on the inverse transformation $h^{-1}(\cdot)$ from the subject to the template. This term is used to constrain the inverse consistency of the transformation $h(\cdot)$ (Christensen, 1999; Christensen and Johnson, 2001; Johnson and Christensen, 2001, 2002). Specifically, these two terms together favor transformations that yield mutually compatible results when deforming the template to the subject and conversely.

The third term is a smoothing constraint for the displacement field, defined on the template space V_T . The selection of the smoothing operators is dependent on the application at hand. In the 3D version (Shen and Davatzikos, 2002, 2003), we used the Laplacian operator. In the 4D formulation described herein, we penalize the displacement fields $u(\mathbf{x}, t)$ that differ significantly from the average displacement field in their 4D neighborhood. For convenience, we separate this smoothness requirement into two components: a spatial smoothing component, E_S^{Spatial} , and a temporal smoothing component, E_S^{Temporal} . α and β are two weighting parameters. In the spatial smoothing component, we use a Gaussian filter g_1 to obtain the average displacement in a 3D spatial neighborhood $n(\mathbf{x})$. The standard deviation in Gaussian filter g_1 changes adaptively as the deformation proceeds. In the temporal smoothing component, we use another Gaussian filter g_2 to obtain the average displacement in a 1D temporal neighborhood $n(t)$. Selecting different sizes of temporal neighborhood and standard deviation for g_2 will lead to a different degree of smoothness on the displacement fields along the temporal dimension. In Results, we will discuss how to select these parameters for accurate longitudinal measurement of brain changes.

Reconstructing the 4D image and the temporal neighborhood structure

Reconstruction of a 4D image by rigidly aligning 3D images

To apply the 4-dimensional warping procedure, we need to reconstruct 4D images from respective 3D images of different time-points. Since the positions and orientations of images from different time-points are not the same, a rigid transformation is first estimated to align the 3D images. The purpose of this step is *not* to capture anatomical changes that occurred in the respective time-interval; this is achieved via 4D warping. Let $\{I_t(\mathbf{x}), t \in [1, N]\}$ be the 3D images before alignment, and $I'_t(\mathbf{x})$ be the 3D images after alignment. We take the first-time image $I_1(\mathbf{x})$ as a standard image, that is, $I'_1(\mathbf{x}) = I_1(\mathbf{x})$. Then, our goal is to rigidly transform the time t image $I_t(\mathbf{x})$ to its previous time image $I'_{t-1}(\mathbf{x})$, where $2 \leq t \leq N$. Many methods are available to estimate the rigid transformation parameters between two 3D images. We employ 3D-HAMMER (Shen and Davatzikos, 2002, 2003) to establish correspondences between two images, $I_t(\mathbf{x})$ and $I'_{t-1}(\mathbf{x})$, and then calculate the rigid transformation parameters from the established correspondences by a least-square estimation

method provided in Umeyama (1991). Let us assume the rigid transformation from $I_t(\mathbf{x})$ to $I'_{t-1}(\mathbf{x})$ is $R_{t \rightarrow t-1}^{Intra}$, where $2 \leq t \leq N$ and $R_{t \rightarrow t-1}^{Intra}$ is a 4×4 matrix with the fourth column as translations. Using this estimated transformation, we can rapidly transform $I_t(\mathbf{x})$ to be $I'_t(\mathbf{x})$. This process is repeated until the last-time image $I_N(\mathbf{x})$ has been aligned as $I'_N(\mathbf{x})$. After rigid alignment, a 4D image is reconstructed from the 3D images. This alignment approach can be used to reconstruct 4D images of both template and subject.

Constructing a 4D template

The 4D template is constructed in a slightly different way from the subject images. In particular, the 3D image of one time-point is repeated, to construct a 4D image. The reason for this difference is that the template is used as a measurement unit, relative to which longitudinal changes of subject images are measured. If the template displayed longitudinal changes, as well, then one would need to know exactly what these changes are and factor them out when measuring an individual’s brain changes. Although this strategy is certainly possible, it would complicate the analysis and display of results, since a changing template would be compared to a changing individual. Therefore, we chose to keep the template fixed temporally, and directly measure an individual’s longitudinal changes. However, a longitudinally changing template might be a preferred strategy if significant longitudinal changes were to be measured, which could render registration difficult unless these changes are present in the template. This has not been the case in our studies.

Transformation of a 4D subject image to the template space

To remove inter-individual size and pose differences, before warping the 4D images, we linearly transform the 4D subject image into the space of the 4D template. This step facilitates subsequent non-rigid image registration. To find this affine transformation, we use the regular moments of the first-time images ($I_1'(\mathbf{x})$) of template and subject to calculate their respective eigenvectors and eigenvalues, and then to estimate a relative 3D linear transformation between template and subject, as we did in the 3D version of HAMMER algorithm (Shen and Davatzikos, 2002, 2003). We apply the estimated 3D transformation to all 3D images in the subject, $\{I'_t(\mathbf{x}) | 1 \leq t \leq N\}$. In this way, the 4D subject image is transformed to the space of the 4D template, thereby producing a transformed 4D subject image. Denote this 3D linear transformation by a 4×4 matrix R^{Inter} , with the fourth column as translations. That means that the original 3D subject image $I_t(\mathbf{x})$ has been transformed to the 4D template space by a total transformation $R^{Inter} R_{t \rightarrow t-1}^{Intra}$ for $2 \leq t \leq N$, and by R^{Inter} for $t = 1$. To avoid interpolation errors when concatenating multiple transformations, we directly transform the original 3D images of the subject, $I_t(\mathbf{x})$, to the space of the 4D template using the combined transformations $R^{Inter} R_{t \rightarrow t-1}^{Intra}$ for the time $2 \leq t \leq N$, and using R^{Inter} for the first time $t = 1$. For convenience, we use $S(\mathbf{x}, t)$ to denote the transformed 4D subject image in the 4D template space, and use $S_t(\mathbf{x})$ to represent a 3D image of the transformed 4D subject at time t .

Rebuilding a temporal neighborhood structure for the 4D subject image

Successive images, such as $S_t(\mathbf{x})$ and $S_{t+1}(\mathbf{x})$, are not necessarily continuous and smooth along the temporal dimension, as

shown schematically in Fig. 1. This is due to longitudinal changes that may occur between two consecutive scans, but also this is due to small misalignment errors. For example, according to the 4D coordinate system, the point A_{t+1} in the image $S_{t+1}(\mathbf{x})$ appears to be an immediate temporal neighbor of the point A_t in the image $S_t(\mathbf{x})$. However, the actual corresponding neighbor of point A_t , in the image $S_{t+1}(\mathbf{x})$, is the point B . Importantly, if we fail to determine the correct neighboring points that are required in Eq. (1), the performance of our approach will be significantly diminished. For example, the total attribute difference/similarity will be integrated from points that are actually not in the neighborhood, and the deformation fields will be smoothed incorrectly along the temporal dimension. To overcome this problem, we use 3D-HAMMER (Shen and Davatzikos, 2002, 2003) to first establish the correspondences between the neighboring 3D images in the 4D subject image $S(\mathbf{x}, t)$, and then use these correspondences to rebuild a new temporal neighborhood structure, which will be used in the procedure of minimizing the cost function in Eq. (1).

In the 4D subject image, $S(\mathbf{x}, t)$, the positions of the corresponding points are not always smooth along the temporal dimension (Fig. 1), as indicated above. However, this is not the case in the 4D template, $T(\mathbf{x}, t)$, since the template was constructed by repeating the “Time 1” image, as described earlier. This poses a very important difficulty. Specifically, one cannot apply customary smoothness constraints in the subject’s 4D space, because such a constraint would ignore the inevitable presence of misalignment and would smooth out important morphological characteristics of the individual. Accordingly, we impose smoothness constraints in the template space only. Specifically, if we follow the backward transformation, that is, the transformation from the 4D subject to the 4D template, we expect to have absolute agreement in the temporal dimension, as shown in Fig. 2. This requirement is part




Fig. 1. Schematic illustration of the procedure used to establish a temporal neighborhood structure for the subject image (only one spatial dimension is shown here, for simplicity). From 4D coordinate system, the point A_{t+1} in the image $S_{t+1}(\mathbf{x})$ appears to be an immediate neighbor of the point A_t in the image $S_t(\mathbf{x})$. However, the correspondence derived from this relationship may be incorrect due to longitudinal changes that may occur between two consecutive scans, and also due to small misalignment errors. For the point A_t , its actual corresponding neighbor in the image $S_{t+1}(\mathbf{x})$ is the point B .




Fig. 2. Schematic diagram explaining the issue of temporal smoothness. No smoothness constraints are imposed on the forward transformation, since inevitable misalignment of the subject images in the temporal dimension is manifested as high-frequency noise, which if removed would smooth out fine morphological details. However, if we follow the backward transformation, we should obtain a temporally smooth warping field. This requirement is built into our optimization and matching scheme. As shown as black dots in the figure, the positions of the corresponding points in the 4D template are very smooth along the temporal dimension. However, the positions of corresponding points in the 4D subject are not always continuous and smooth along the temporal dimension, as indicated by black circles. Here, the point $C_\tau(\mathbf{x}, t)$ in the image of time τ is the temporal correspondence of the point (\mathbf{x}, t) in the image of time t .

of our optimization and matching scheme, which can be mathematically formulated as:

$$E_S^{\text{Temporal}} = \sum_{\mathbf{x} \in V_S} \sum_{t=1}^N \left(h^{-1}(\mathbf{x}, t) - \frac{\sum_{\tau \in n(t)} g_2(\tau - t) h^{-1}(C_\tau(\mathbf{x}, t), \tau)}{\sum_{\tau \in n(t)} g_2(\tau - t)} \right)^2, \quad (7)$$

where $h^{-1}(\mathbf{x}, t)$ is the backward transformation on the subject point (\mathbf{x}, t) . $C_\tau(\mathbf{x}, t)$ denotes the temporal correspondence of the subject point (\mathbf{x}, t) in the image of time τ , $C_\tau(\mathbf{x}, t)$. The definition of this temporal smoothness constraint is similar to that in Eq. (6). In our study, we used Eq. (7) to constrain the temporal smoothness of the transformation.

Attribute-based deformable matching

In this section, we describe the 4D deformation strategy of our approach. Generally speaking, the 4D template is warped to the subject by attempting to match the attribute vectors of a number of points, which we call active points. These are the points with relatively distinctive attribute vectors, which can be identified relatively more reliably based on their morphological signatures, compared to other more ambiguous points. The selection of the active points is determined purely from their attribute vector values in a training set of images. To evaluate a candidate match, a subvolume (i.e., the tissue in a neighborhood around the point of interest) is defined and warped toward the subject, for assessing the goodness of fit of the attribute vectors in that subvolume. More and more active points are added hierarchically to the search strategy. The details of the method are given below.

Hierarchical selection of the active points

The deformation of the template $T(\mathbf{x}, t)$ to the subject $S(\mathbf{x}, t)$ is driven by a hierarchically selected set of active points. This strategy is aimed at speeding up performance and reducing local minima, which in part result from ambiguities in determining the matching pairs of points. Some structures have very distinctive attribute vectors, for example, the roots of sulci and the crowns of the gyri, which can be identified more reliably than other cortical points. Therefore, in the deformable matching procedure, we first evaluate candidate matches of points with distinctive attribute vectors, and then gradually consider points with less distinctive attribute vectors. Effectively, this procedure approximates what would be a very high-dimensional (equal to the number of points in 4D image) cost function in Eq. (1), by a significantly lower dimensional function of only the active points. This latter function has few local minima, because it is a function of the coordinates of points for which relatively unambiguous matches can be found.

In terms of practical implementation of this approach, we first rank the points of 4D images according to the distinctiveness of their attribute vectors. We then select the top-ranked points as the active points to drive the initial deformation of the template to the subject. As the deformation progresses, we gradually add other lower-ranked points to the set of the active points, for further deforming the template to the subject. In the final stage, we use all points in the template as the active points to drive the deformations.

Inverse consistent transformation

An important characteristic of this matching process is that it is inverse consistent. A transformation is inverse consistent if it finds consistent results when matching a template to a subject versus when it matches the subject to the template (Christensen, 1999; Christensen and Johnson, 2001; Johnson and Christensen, 2001, 2002; Shen and Davatzikos, 2002, 2003). However, most automated image warping algorithms are not inverse consistent. An approach to enforcing inverse consistency in a warping process was described in Christensen (1999), Christensen and Johnson (2001) and Johnson and Christensen (2001, 2002). However, that approach is computationally very demanding. In our approach, we enforce inverse consistency only on the active points, since the warping transformation elsewhere is determined via interpolation from the active points. Our methodology for obtaining inverse consistent matches on the active points was described in detail in Shen and Davatzikos (2002, 2003).

Subvolume deformation strategy

As we mentioned earlier, to increase robustness to potentially false matches of active points, we evaluate the similarity of attribute vectors in the entire subvolume around an active point, not simply the similarity of active points. In particular, for each template active point, we search for several target points with similar attribute vectors in its neighborhood. For each candidate match, the similarity of attribute vectors in the respective 4D subvolumes is evaluated. This is an important aspect of our method. Say, for example, that a false match within an active point's neighborhood happens to have similar attribute vector. It is unlikely that all points in the spatial as well as temporal neighborhood also have matching attribute vectors, unless this is a true match. The point in the subject 4D image that presents the highest

subvolume-derived attribute vector similarity, is defined as a correspondence, provided that the similarity is above a pre-specified threshold. This prevents active points that do not find good matches from actively deforming the volume, as would be the case of a single sulcus in the template deforming to a double sulcus in the subject brain. These multiple levels of evaluation of goodness of correspondence add robustness that is necessary in view of the complexity of the human brain anatomy. On the other hand, correspondences in the reverse direction, that is, from the subject to the template, are also examined. If these correspondences turn out to be inconsistent with the forward process described above, then a compromise is reached so that inverse consistency is satisfied.

This process is applied in a hierarchical way. In particular, the size of the subvolume is initially very large, since at that stage there are only a few template active points distributed sparsely in the 4D space. It decreases gradually as the deformation progresses, and it reduces to include only immediate neighbors in the final stages. In a 4D image, the range of temporal dimension is usually much smaller than the ranges of the spatial dimensions (such as $256 \times 256 \times 124$). Accordingly, the 4D subvolume is designed to be flat in the temporal dimension. The size of the Gaussian interpolation filters used to smear the displacements of the active points to the remainder of the 4D volume are also designed to get progressively smaller, as the density of active points increases.

Summary of the algorithm

Our 4D algorithm is implemented in a multi-resolution way analogous to the 3D formulation that was described in Shen and Davatzikos (2002). For completeness, we include the entire description here.

1. Rigidly align the 3D images of the subject acquired at different time-points, $\{I_t(\mathbf{x}), 1 \leq t \leq N\}$. We use 3D-HAMMER to establish the correspondences between the 3D images of the neighboring time-points, and then calculate the rigid transformation parameters ($R_{t \rightarrow t-1}^{Intra}$) from the established correspondences (Umeyama, 1991) and align the time t image to its previous image. The aligned 3D images are represented as $\{I'_t(\mathbf{x}), 1 \leq t \leq N\}$.

2. Linearly transform the 4D subject to the space of the 4D template, using the geometric moments of the first-time images of the subject and the template (Shen and Ip, 1997). The calculated transformation (R^{Inter}) will bring 4D subject to the space of 4D template. To minimize the effect of digitization and interpolation errors, we combine the linear transforms $R_{t \rightarrow t-1}^{Intra}$ and R^{Inter} , and directly transform the original 3D subject images $I_t(\mathbf{x})$ to the space of 4D template by using a combined transformation $R^{Inter} R_{t \rightarrow t-1}^{Intra}$. The final linearly transformed 4D subject image is represented as $S(\mathbf{x}, t)$, where $1 \leq t \leq N$.

3. Compute the attribute vectors for all points in the 4D subject and template images. Since brain tissues have been labeled to GM, WM and CSF, we can calculate the respective attributes such as GMIs from the local structures of GM, WM and CSF for each voxel.

4. Determine a small set of active points with the most distinctive attribute vectors from the 4D subject, $S(\mathbf{x}, t)$.

5. Hierarchically select a set of active points in the 4D template, $T(\mathbf{x}, t)$. The number of active points is small initially, and it gradually increases as the deformation progresses. Since

bigger structures tend to dominate over smaller structures, simply because they contribute more active points, we have implemented a spatially adaptive mechanism for selecting proportionately more active points from small structures.

6. For each active point in the subject image, search for its corresponding point in the template image of the same time-point; no temporal warping is allowed, since the time-points of image acquisitions are known in advance.

7. For each template active point, search for the best match in the target image. If the correspondence is found, then a 4D subvolume (spatial and temporal) around that active point will be deformed; otherwise, that active point will not produce any deformation to the template in the current stage. The size of 4D subvolume decreases with the progress of the deformation; in the final stages, only the immediate neighbors are included.

8. Smooth the deformation fields using Gaussian kernels of progressively smaller width, and using local elastic constraints.

9. Repeat steps 5–8 until the algorithm converges.

Results

4D-HAMMER is designed to accurately measure the longitudinal changes, particularly in small structures. This is especially required in our studies seeking early and subtle markers of brain disease. Usually, manual raters are extensively trained to label structures such as hippocampus. However, manual segmentation requires extensive human interaction and considerable training. Most importantly, both intra-rater reproducibility and inter-rater agreement are difficult to achieve in a longitudinal study, particularly when small changes are to be measured. 4D-HAMMER could prove particularly useful in measuring the small longitudinal changes. As part of our evaluation procedure, we will compare the performance of 4D-HAMMER with that of experienced rater, by using small structure such as hippocampus as a testing example.

Two sets of experiments are described in this section, for demonstrating the performance of this 4D atlas matching approach in estimating longitudinal changes of hippocampal volumes. The hippocampal regions in the template have been labeled via manual definition of a single template brain. Then, by warping the template to the subject using the proposed approach, we can simultaneously warp the mask of the template hippocampi to each subject, thereby obtaining the segmentation of hippocampi from each subject and each time-point. The first set of experiments uses simulation data, in which an atrophy is simulated in part of the hippocampus. The second set of experiments uses images from the Baltimore Longitudinal Study of Aging (BLSA) (Resnick et al., 2000), and compares automated segmentation results with results obtained by well-trained experts. In both sets of experiments, the performances of 4D- and 3D-HAMMER are compared, which reveals that 4D-HAMMER performs much better than 3D-HAMMER in estimating longitudinal changes, particularly for small structures such as the hippocampus. This is because 4D warping algorithm simultaneously establishes the longitudinal correspondences in the individual, and also the inter-subject correspondences between the template and the individual. The performance of 4D-HAMMER in correctly establishing correspondences is visually demonstrated by Fig. 3.

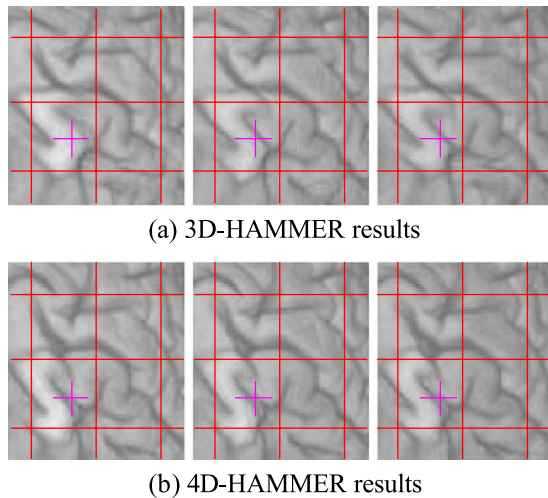


Fig. 3. Comparison of 3D-HAMMER and 4D-HAMMER in warping three temporal MR images of a selected subject to the template space. Since the temporal images of the template are identical, the warped subject images of different time-points should be very similar. In 3D-HAMMER, different time-point images were warped independently, thereby leading to unavoidable inconsistencies in image matching among different time-points, as indicated by the pink crosses in (a). On the contrary, the warping results by 4D-HAMMER are very consistent, as reflected by the almost identical warped images in (b). The pink crosses, with the same 3D coordinates respectively in (a) and (b), are used as references for comparing the corresponding warpings.

Simulation data

Three subjects were selected from the BLSA project, and used to simulate atrophy on their hippocampi. For each subject, the five simulated MR images were created by the following way. First, we used the original image as a first-time image ($t = 1$). We then used a biomechanical model to introduce a contracting stress




Fig. 4. Simulating atrophy in one part of each hippocampus. (a) Left and right hippocampi of a subject. About 5% of atrophy was simulated in the yellow parts of the hippocampi in (a), which resulted in the shrunken hippocampi [displayed as yellow in (b)]. Notice that the non-overlapping part of the original and shrunken hippocampi is shown as green in (b), and it corresponds to about 5% atrophy simulated.

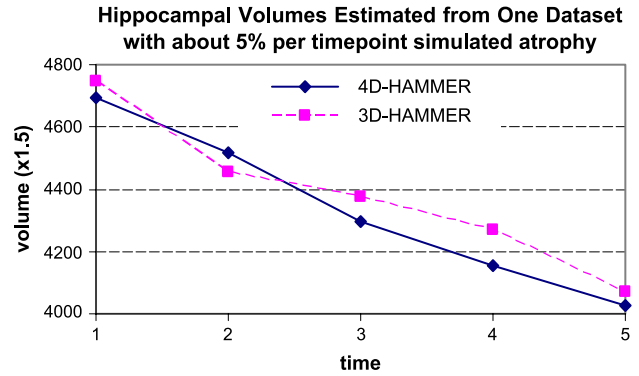


Fig. 5. Comparing the performances of 4D- and 3D-HAMMER in segmenting hippocampi from the simulated MR images of a typical subject. The pink dotted curve represents the hippocampus segmentation results by 3D-HAMMER, and the dark-blue curve denotes the results obtained by 4D-HAMMER. The result of 4D-HAMMER is smoother, and it provides a better estimate of the true underlying atrophy.

within part of the hippocampus, which shrunk the hippocampi of the current-time image ($t = 1$) by about 5%, yielding a next-time image ($t = 2$). By repeating this procedure again and again, we finally obtained five simulated images whose hippocampal volumes shrunk at a rate of about 5% per time-point along the temporal dimension. Fig. 4 gives an example of this hippocampus shrinking procedure. The yellow part in Fig. 4(a) indicates the selected hippocampal region on which the atrophy was simulated. Fig. 4(b) shows an example of the shrunken hippocampi, represented by the yellow part; the green colors denote the lost hippocampal volume. By applying both 4D- and 3D-HAMMER on the simulated images of three subjects, we could compare the performances of these two methods in estimating the shrinkage rate of hippocampal volume. Fig. 5 gives a typical result on one subject, where the pink dashed curve denotes the hippocampal volumes detected by 3D-HAMMER, while the dark-blue curve denotes the result obtained by 4D-HAMMER. Since an almost constant rate of atrophy was simulated along the temporal direction, we expect that the estimated hippocampal volumes are distributed at a nearly straight line. Fig. 5 shows that 4D-HAMMER produces better results than 3D-HAMMER. Moreover, we can compare the average hippocampal volume of those three

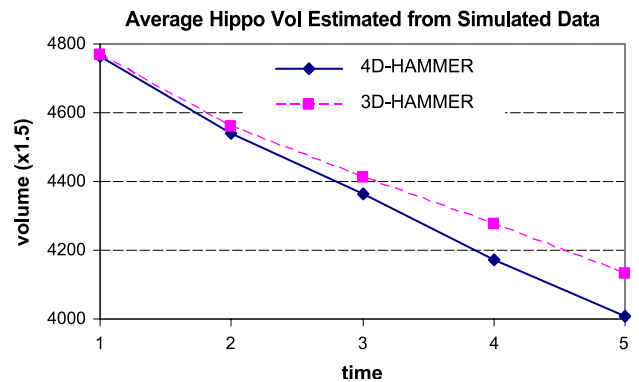


Fig. 6. Average hippocampal volumes estimated from simulated images of three subjects, by 4D- and 3D-HAMMER. The total atrophy was 16%. Estimated atrophy: 15.9% by 4D-HAMMER and 13.3% by 3D-HAMMER.

subjects, estimated respectively by 4D- and 3D-HAMMER. The average percentage of atrophy that we simulated on the hippocampi of three subjects is 16.0% totally from time 1 to time 5. (Notably, if the shrinkage rates was exactly 5% per time-point along the temporal dimension, then the total shrinkage percentage would be $(1 - (1 - 0.05)^4) \times 100\% = 18.5\%$; numerical errors in the shrinking process resulted in simulated atrophy of a slightly smaller 16%.) The 4D-HAMMER estimates 15.9% atrophy, while the 3D-HAMMER estimates only 13.3% atrophy. Fig. 6 displays these results.

Real data

4D-HAMMER was also used to estimate the longitudinal changes of hippocampal volumes in the real MR images of nine elderly subjects aged 59–78, selected from our ongoing studies in the BLSA (Resnick et al., 2000). The reason of selecting only nine subjects for demonstrating the performance of 4D-HAMMER is because we only have manual segmentation results for those nine subjects. Actually, there are over 100 subjects in BLSA project, for which annual MR images are available for each subject and have been collected over a 9-year period. In this paper, the images from the first 5 years were used. In an earlier publication (Shen et al., 2002), we described the process of manually segmenting the hippocampus, to compare automated methods with manual ratings.

Fig. 7 shows the hippocampal volumes of a typical subject, detected by 4D- and 3D-HAMMER, respectively. It can be observed that the results from 4D-HAMMER are much smoother than those from 3D-HAMMER. Fig. 8 shows the average hippocampal volumes of nine subjects estimated by 3D- and 4D-HAMMER, respectively. 3D-HAMMER resulted in noisy longitudinal estimations, while 4D-HAMMER produced smooth longitudinal estimations. Although the results by 3D-HAMMER can be smoothed, the longitudinal change still deviates significantly from the one estimated by an experienced rater. For example, the percentage of hippocampal shrinkage detected during the 4 years (from year 1 to 5) is only 2.1% by 3D-HAMMER, compared to 5.5% by 4D-HAMMER, and 5.7% using manual definition by the experienced rater. A scatterplot of the measurements of longitudinal changes determined by the experienced rater and by 4D-HAMMER is also shown in Fig. 9(a). It can be observed that most of these nine points are distributed around the diagonal line, indicating reason-

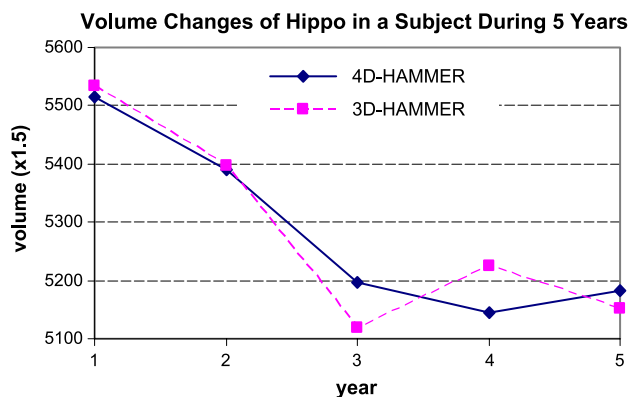


Fig. 7. Comparing 4D- and 3D-HAMMER in estimating the longitudinal changes of hippocampi from a real subject.

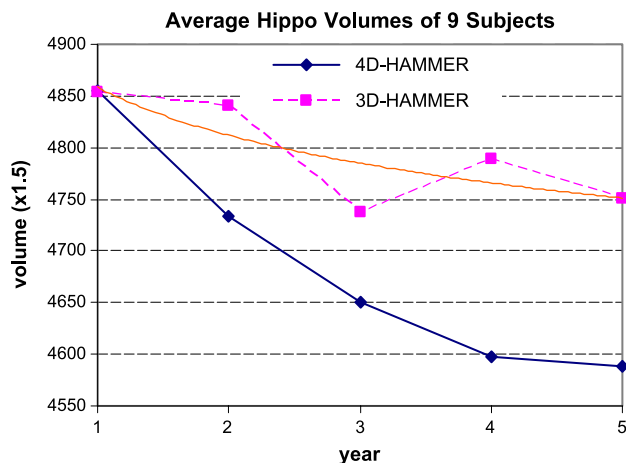


Fig. 8. Average hippocampal volumes of the nine subjects, estimated by 4D- and 3D-HAMMER. The shrinking percentage estimated by 4D-HAMMER is very close to that achieved by an experienced rater. Additionally, the results by 4D-HAMMER are smoother than those by 3D-HAMMER. Although the results by 3D-HAMMER can be smoothed as displayed by orange solid curve, the estimated longitudinal change still deviates significantly from the one obtained by the experienced rater.

able correlation in view of the unavoidable errors in calculating longitudinal differences. There is one case where 4D-HAMMER measured 0% change and the expert showed 20% change. This is because the low-intensity voxels inside of hippocampus were labeled as non-hippocampal tissue by the expert, while they were labeled as hippocampal tissue by 4D-HAMMER, since solid hippocampus template was used in our study. It is worth noting that the manual rater produced the negative shrinkage percentages for the three subjects, which is indeed biologically unlikely for the elderly subjects examined in our study. For comparison, the absolute hippocampal volumes of nine year 1 images determined by our method and manual rater are provided in Fig. 9(b). The correlation coefficient between our automated method and the manually defined gold standard is 0.89. In general, the performance of 4D-HAMMER is statistically close to that of an experienced rater, and also 4D-HAMMER is more robust and accurate in estimating longitudinal changes, compared to 3D-HAMMER.

As we indicated in Mathematical optimization subsection, selecting different sizes of temporal neighborhood and standard deviation for Gaussian filter g_2 will lead to different amounts of smoothness on the displacement fields and further on the longitudinal measurements. Certainly, more smoothing on temporal parameters results in smoother longitudinal measurements. However, the goal of 4D-HAMMER is not only to make the longitudinal measurement smooth, but also to make the algorithm-based estimation close to the actual amount of longitudinal changes. Therefore, the temporal smoothing parameters should be determined by this criterion. For convenience, let N_t denote the size of temporal neighborhood $n(t)$, used for Gaussian filter g_2 . $N_t = 1$ means that only immediate temporal neighbors are used to smooth the displacement field along the temporal dimension. Let σ denote the standard deviation of Gaussian filter g_2 . Three curves in Fig. 10 represent the average hippocampal volumes of the nine subjects estimated by 4D-HAMMER, respectively using three different sets of temporal smoothing

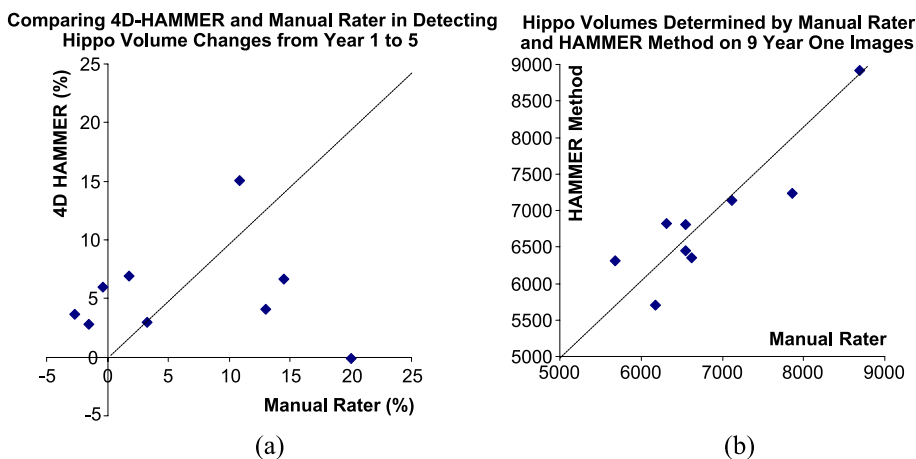


Fig. 9. (a) Visual demonstration for the correlation between longitudinal changes determined by 4D-HAMMER and by an experienced rater. [Procedures for obtaining manual ratings were described in Shen et al. (2002).] Given variations among raters themselves in measuring longitudinal change of such a small magnitude in structures such as the hippocampus, a reasonable correlation is apparent. It is worth noting that the manual ratings produced negative shrinking percentages for three subjects, which is actually biologically implausible for the elderly subjects examined in our study. (b) A scatterplot of absolute hippocampal volumes determined by our automated method and the experienced rater, which indicates a good agreement between automated and manual methods (correlation coefficient 0.89).

parameters. The parameters (2.0,1) in the figure means $\sigma = 2.0$ and $N_t = 1$. It can be observed from the figure that the different sets of temporal smoothing parameters can lead to different estimation results, especially if the temporal neighborhood is very small, as should be expected. The use of smoothing parameters (5.0,5) produced the smoothest result, as well as the best result according to the manual segmentation by an experienced rater. Therefore, we selected $\sigma = 5.0$ and $N_t = 5$ as the temporal smoothing parameters for our study. Notably, the dark-blue curve in Fig. 8 was obtained

by using the temporal smoothing parameters (5.0,5), and it is identical to the dark-blue curve in Fig. 10.

Conclusion

We have developed a robust and accurate 4D atlas matching method for registration and warping of 4D images, and for estimating the longitudinal morphological changes of brain structures. The experimental results using both simulated and real data show that 4D-HAMMER can provide smooth and accurate longitudinal estimations even for small difficult structures such as hippocampus. The main characteristics of our approach are summarized next. First, an attribute vector is defined as a morphological signature for every point in a 4D image. The attribute vector is developed to be as distinctive as possible, thereby reducing the ambiguity in correspondence matching and finally avoiding local minima. Second, according to the distinctiveness of the attribute vectors, the points in the 4D space are successively selected as active points for hierarchically deforming the template to the subject. This helps us avoid local minima and also speed up the algorithm. Third, our approach is made further robust to local minima and noise, by using a 4D subvolume matching and deformation strategy. Finally, the displacement fields are preserved to be smooth at both spatial and temporal domains, via appropriate constraints defined in both domains.

Some extensions of the 4D-HAMMER algorithm are possible. Currently, we use the tissue-segmented images to calculate the local morphological signature for each voxel of a sequence of images. Since the images of different time-points are segmented independently in our current study, the tissue segmentations along the temporal dimension may be inconsistent. If all temporal MR images of an individual can be considered simultaneously in the process of tissue segmentation, then the temporal consistency of the tissue segmentations can be increased. When the new tissue segmentation method is available, it will be used in our 4D-HAMMER algorithm. Also, for further testing the performance

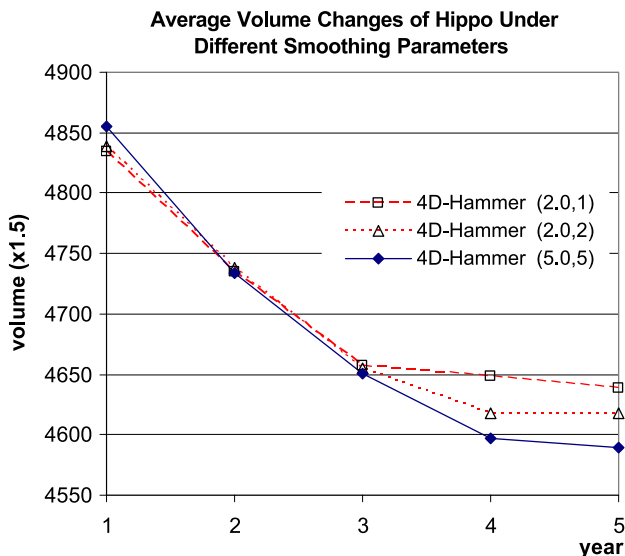


Fig. 10. The estimated hippocampal volumes varying with the smoothing parameters used in the Gaussian filter g_2 . The parameters (2.0,1) means that the size of the temporal neighborhood is 1 ($N_t = 1$) and the standard deviation of the Gaussian filter g_2 is 2.0 ($\sigma = 2.0$). The use of smoothing parameters (5.0,5) produced the smoothest result, and actually the best result according to the manual segmentation achieved by an experienced rater.

of 4D-HAMMER algorithm, we are planning to apply 4D-HAMMER to over 100 subjects and multiple brain regions in the BLSA neuroimaging study, for which MR images have been collected over a 9-year period.

Acknowledgments

We thank Tianming Liu for preparing the simulation data and Susan Resnick from the Baltimore Longitudinal Study of Aging for providing the data sets.

References

- Bajcsy, R., Lieberman, R., Reivich, M., 1983. A computerized system for the elastic matching of deformed radiographic images to idealized atlas images. *J. Comput. Assist. Tomogr.* 7 (4), 618–625.
- Bookstein, F.L., 1989. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (6), 567–585.
- Brejl, M., Sonka, M., 2000. Object localization and border detection criteria design in edge-based image segmentation: automated learning from examples. *IEEE Trans. Med. Imag.* 19, 973–985.
- Christensen, G.E., 1999. Consistent linear-elastic transformations for image matching. *Information Processing in Medical Imaging, LCNS 1613*. Springer-Verlag, Visegrád, Hungary, pp. 224–237.
- Christensen, G.E., Johnson, H.J., 2001. Consistent image registration. *IEEE TMI* 20 (7), 568–582.
- Christensen, G.E., Rabbitt, R.D., Miller, M.I., 1996. Deformable templates using large deformation kinematics. *IEEE Trans. Image Process.* 5 (9), 1435–1447 (September).
- Chui, H., Win, L., Schultz, R., Duncan, J., Rangarajan, A., 2001. A unified feature registration method for brain mapping. *Information Processing in Medical Imaging, Davis, CA, USA, June 18–22*, pp. 300–314.
- Chung, M.K., Worsley, K.J., Paus, T., Cherif, C., Collins, D.L., Giedd, J.N., Rapoport, J.L., Evans, A.C., 2001. A unified statistical approach to deformation-based morphometry. *NeuroImage* 14 (3), 595–606 (September).
- Davatzikos, C., Bryan, R.N., 1996. Using a deformable surface model to obtain a shape representation of the cortex. *IEEE Trans. Med. Imag.* 15, 785–795 (Dec.).
- Evans, A.C., Dai, W., Collins, L., Neeling, P., Marett, S., 1991. Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. *SPIE Proc., Image Process.* 1445, 236–246.
- Fox, P.T., Mintum, M.A., Reiman, E.M., Reichle, M.E., 1988. Enhanced detection of focal brain responses using inter-subject averaging and distribution analysis of subtracted PET images. *J. Cereb. Flow Metab.* 8, 642–653.
- Freeborough, P.A., Fox, N.C., 1998. Modeling brain deformations in Alzheimer's disease by fluid registration of serial 3D MR images. *J. Comput. Assist. Tomogr.* 22, 838–843.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.P., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.*, 189–210.
- Gee, J.C., Barillot, C., Briquer, L.L., Haynor, D.R., Bajcsy, R., 1994. Matching structural images of the human brain using statistical and geometrical image features. *Proc. SPIE Vis. Biomed. Comput.* 2359, 191–204.
- Johnson, H.J., Christensen, G.E., 2001. Landmark and intensity-based, consistent thin-plate spline image registration. *Information Processing in Medical Imaging, Davis, CA, USA, June 18–22*, pp. 329–343.
- Johnson, H.J., Christensen, G.E., 2002. Consistent landmark and intensity-based image registration. *IEEE TMI* 21 (5), 450–461.
- Joshi, S.C., Miller, M.I., Christensen, G.E., Banerjee, A., Coogan, T., Grenander, U., 1995. Hierarchical brain mapping via a generalized Dirichlet solution for mapping brain manifolds. *Proc. SPIE Conf. Geom. Methods Appl. Imag.* 2573, 278–289 (July).
- Lo, C.H., Don, H.S., 1989. 3-D moment forms: their construction and application to object identification and positioning. *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (10), 1053–1064 (Oct.).
- Miller, M.I., Christensen, G.E., Amit, Y., Grenander, U., 1993. Mathematical textbook of deformable neuroanatomies. *Proc. Natl. Acad. Sci. U. S. A.* 90, 11944–11948.
- Resnick, S.M., Goldszal, A.F., Davatzikos, C., Golski, S., Kraut, M.A., Metter, E.J., Bryan, R.N., Zonderman, A.B., 2000. One-year age changes in MRI brain volumes in older adults. *Cereb. Cortex* 10, 464–472.
- Rohr, K., 1999. Image registration based on thin plate splines and local estimates of anisotropic landmark localization uncertainties. *Lect. Notes Comput. Sci.: MICCAI'98* 1496, 1174–1183.
- Roland, P.E., Graufelds, C.J., Wahlin, J., et al., 1994. Human brain atlas: for high-resolution functional and anatomical mapping. *Hum. Brain Mapp.* 1, 173–184.
- Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imag.* 18 (8), 712–721 (August).
- Shen, D., Ip, H.H.S., 1997. Generalized affine invariant image normalization. *IEEE Trans. Pattern Anal. Mach. Intell.* 19 (5), 431–440 (May).
- Shen, D., Davatzikos, C., 2002. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imag.* 21 (11), 1421–1439 (Nov.).
- Shen, D., Davatzikos, C., 2003. Very high resolution morphometry using mass-preserving deformations and HAMMER elastic registration. *NeuroImage* 18 (1), 28–41 (Jan.).
- Shen, D., Moffat, S., Resnick, S.M., Davatzikos, C., 2002. Measuring size and shape of the hippocampus in MR images using a deformable shape model. *NeuroImage* 15 (2), 422–434 (Feb.).
- Thirion, J.P., 1996. Non-rigid matching using demons. *Proc. IEEE Conf. Comp. Vis. Pattern Recognit.*, 245–251.
- Thirion, J.P., Monga, O., Benayoun, S., Guezic, A., Ayache, N., 1992. Automatic registration of 3-D images using surface curvature. *SPIE Proc., Math. Methods Med. Imag.* 1768, 206–216.
- Thompson, P., Toga, A.W., 1996. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. Med. Imag.* 15, 402–417.
- Umeyama, S., 1991. Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. PAMI* 13 (4), 376–380.
- Vaillant, M., Davatzikos, C., 1999. Hierarchical matching of cortical features for deformable brain image registration. *Lect. Notes Comput. Sci.: Information Processing in Medical Imaging* 1613, 182–195 (June).
- Weese, J., Penney, G.P., Desmedt, P., Buzug, T.M., Hill, D.L.G., Hawkes, D.J., 1997. Voxel-based 2-D/3-D registration of fluoroscopy images and CT scans for image-guided surgery. *IEEE Trans. Inf. Technol. Biomed.* 1 (4), 284–293.
- Xue, Z., Shen, D., Davatzikos, C., submitted for publication. Determining correspondence in 3D MR brain images using attribute vectors as morphological signatures of voxels. *IEEE Trans. Med. Imag.*