

Very High Resolution Morphometry Using Mass-Preserving Deformations and HAMMER Elastic Registration

Dinggang Shen *and* Christos Davatzikos

Section of Biomedical Image Analysis
Department of Radiology, University of Pennsylvania, Philadelphia, PA

Abstract

This paper presents a very high resolution voxel-based morphometric method, by using a mass-preserving deformation mechanism and a fully automated spatial normalization approach, referred to as HAMMER. By using a hierarchical attribute-based deformation strategy, HAMMER partly overcomes limitations of several existing spatial normalization methods, and it achieves a level of accuracy that makes possible morphometric measurements of spatial specificity close to the voxel dimensions. The proposed method is validated by a series of experiments, with both simulated and real brain images.

I. Introduction

Volumetric analyses of brain images have traditionally relied on manually defining a number of regions of interest (ROIs), and performing statistical analysis on the volumes of these ROIs. That approach suffers from many well-known drawbacks, including high demand for human effort which limits the number of subjects and ROIs that can be examined, subjectivity and lack of reproducibility, and the need for *a priori* knowledge of regions of interest. The latter is an important limitation, because it is not possible to know in advance which regions might be affected by a disease or differ between two populations. Even if the regions were approximately known, the ROIs from which volumetric measurements are taken are likely to include other surrounding regions, which blurs the results and reduces statistical power.

In order to overcome these limitations, several investigators have pursued voxel-based methods that utilize automatically or semi-automatically determined shape transformation and analysis techniques. These methods generally fall under three different categories. The first category includes methods that measure the spatial transformation that is necessary to deform a template of brain anatomy to each

individual in the study. Several variants of these approaches exist, depending on the analysis of the spatial transformation [Miller, 1993; Davatzikos, 1996, 1998a, 2001a; Thompson, 1997, 2000; Machado, 1998; Gee, 1993; Briquer, 1997; Bookstein, 1989; Joshi, 1997; Subsol, 1998; Freeborough, 1998; Thirion, 1992, 1996; Chui, 2001; Chen, 1999; Wang 2000; Evans, 1991; Rueckert, 1999; Hellier, 1999; Collins, 1999; Christensen, 2001]. These approaches are perhaps the most rigorous way of measuring shape. However, their current limitation is that they rely heavily on a perfect registration between the template and the individuals. Even small registration errors may significantly reduce the accuracy of these methods [Bookstein, 2001]. The second category involves the method commonly referred to as voxel-based morphometry (VBM) [Wright, 1995; Ashburner, 1998]. This approach uses a relatively lower dimensional spatial transformation to remove overall shape differences across individuals, and examines the residual variability of gray and white matter in the normalized space. If there is some localized atrophy, presumably it will be manifested by relatively smaller amounts of brain tissue in the respective location in the normalized space. This technique has been used in several studies during the past few years. Its main limitation is its somewhat heuristic nature, which stems from the fact that VBM measures residual variability after spatial normalization, which does not precisely define any anatomical variable. Indeed, the results of this approach will be different if different registration methods are used. If the registration is perfect, no residual variability, and therefore no local volumetric effects, can be measured. The same is true if the registration is very poor, in which case severe blurring will confound the results.

The method examined in this paper falls under a third category, which includes *tissue-preserving shape transformations*, which were initially introduced in [Goldszal, 1998; Davatzikos, 1998b], and later in [Ashburner, 2000; Davatzikos, 2001b]. These approaches are motivated by the fact that spatial normalization changes the anatomy to be measured, in an effort to place this anatomy into a canonical reference space. In order to account for such changes, these methods take special care to preserve the total amount of tissue of any structure or part of it. This is achieved by changing the density of the tissue within the structure, according to the amount of expansion or contraction that the spatial transformation imposes. A physical analog is the squeezing of a rubber object, which changes the density of the rubber, in order to maintain the same total mass in the object. Regional volumetric measurements are then performed via the resulting tissue density maps. In [Goldszal, 1998; Davatzikos, 2001b], we called this approach Regional Analysis of Volumes Examined in Normalized Space (RAVENS). Similarly, for the case of spatially normalizing functional images (i.e fMRI), the signal preserving transformations could also be applied, in order to retain the total amount of “activation” in different subjects.

Notably, the performance of the voxel-based methods in the third category are also highly related to the accuracy of the image registration. In this paper, we describe a fully automated spatial transformation approach, referred to as Hierarchical Attribute Matching Mechanism for Elastic Registration (HAMMER), which achieves a very high accuracy of registration, thus allowing for morphometric measurements of spatial specificity close to the voxel dimensions. We use HAMMER in conjunction with the mass-preserving RAVENS framework on both simulated and real images, in order to perform regional volumetric measurements.

II. Methods

RAVENS tissue density maps. We now describe briefly the RAVENS methodology that was originally presented in [Goldszal, 1998; Davatzikos, 1998b; Davatzikos, 2001b], and which is based on a volume-preserving spatial transformation. As we mentioned earlier, in order to better describe that method, we draw upon a physical analogy: squashing a rubber object increases the material density of the object. Similarly, stretching the object reduces the material density, since no new material is generated during the object's deformation. Similarly, in the RAVENS framework, if an individual's ventricles are deformed into conformation with a template's ventricles, local expansion or contraction changes the local density of CSF. This is also true for GM and WM structures, as well as for arbitrary subdivisions of them. Since the original information about volumes of brain structures and any arbitrary partitions of them is converted into tissue densities, and since these tissue density RAVENS maps are registered, local differences or changes in volumes can be quantified by respective changes in the RAVENS maps. Fig. 1 demonstrates this principle. A similar approach was also adopted in [Ashburner, 2000].

HAMMER elastic registration. Regardless of the approach used for voxel-based analysis, spatial registration and normalization is a critical step in voxel-based morphometry. So far, many registration methods have been developed [Thirion, 1992, 1996; Subsol, 1998; Gee, 1994; Chui, 2001; Chen, 1999; Pizer, 1999; Evans, 1991; Bajcsy, 1983; Breijl, 2000; Rohr, 1999; Joshi, 1995; Vemuri, 2001; Wang 2000; Stefansic, 2000; Thévenaz, 2000; Rueckert, 1999]. In this section we describe an approach that is based on a Hierarchical Attribute Matching Mechanism for Elastic Registration (HAMMER). Some of the technical details of an earlier implementation of HAMMER can be found in [Shen, 2001; Shen, 2002]. HAMMER was developed to overcome two common limitations of existing fully automated registration methods, which are the following:

Limitation 1. High-dimensional image matching relies on optimization of some matching function, often reflecting the similarity between two images. This function is known to have many local minima, i.e. solutions that appear to be locally optimal but can actually represent poor matches. Local minima are caused not only by the many degrees of freedom, but by the complexity of the underlying brain anatomy, as well.

Limitation 2. Many methods determine the high-dimensional deformation field by attempting to match the intensities of the warped volume with those of the target volume. While this approach is attractive because of its fairly straightforward implementation, it does not guarantee that anatomically meaningful matches are generated, since image similarity does not necessarily imply anatomical correspondence. Moreover, intensity matching contributes to the problems of *Limitation 1*, since it increases ambiguity in the matching procedure, and therefore introduces local minima. For example, consider a point lying in the WM. There are hundreds of thousands of points having similar image intensity. With no additional mechanisms to distinguish between different WM points, ambiguity is created and can result in poor matches and local minima.

HAMMER addresses both issues, as described next. We first describe how we address *Limitation 2*.

Addressing Limitation 2. We have used the concept of an *attribute vector*, which is a collection of geometric attributes defined on every single voxel in a volumetric image and which reflect the geometric properties of the anatomy in the vicinity of that voxel. The attribute vector includes intensity and edge type information, but it also includes a number of *geometric moment invariants* (GMIs), which are quantities that have been used successfully in computer vision for image understanding. GMIs are rotation-invariant quantities calculated at different scales, i.e. for different sizes of the neighborhood around a voxel under consideration, and they reflect the anatomy in the neighborhood of the voxel. Ideally, if a large number of GMIs are used at various scales, the resulting attribute vector will uniquely identify the respective voxel among other voxels of similar intensities, unless the anatomy around that voxel is very similar to the anatomy around another voxel in an image. This is shown in Fig. 2. In practice, we calculate a limited number of GMIs, in order to reduce computational requirements. The attribute vector is a *voxel signature*, which dramatically reduces ambiguity in image matching, although it does not necessarily identify all voxels uniquely.

Addressing Limitation 1. This aspect of HAMMER is presented second, since it relies, in part, on the attribute vectors. In particular, in order to overcome the problem of local minima of the matching function, we use a sequence of hierarchical approximations of this energy function by a number of lower dimensional smooth energy functions. This is achieved by hierarchically selecting the driving

features that have distinctive attribute vectors, thus drastically reducing ambiguity in finding correspondence. For example, in the initial stages of the algorithm, a small number of basis functions, corresponding to the voxels with the distinctive attribute vectors, are used to build up the energy function. These functions are designed very carefully, so that they result in an energy function with very few local minima. Their design is based on voxels with relatively more unique attribute vectors. Specifically, a small set of voxels are initially selected to drive the deformation procedure. These voxels are automatically selected to have rather distinctive attribute vectors, and typically lie on roots of sulci or crowns of gyri, as well as on other distinctive brain structures such as the anterior horn of the ventricles or of the caudate nucleus. A Gaussian kernel is used to propagate the displacement of these voxels to other voxels in their vicinity, as shown in Fig. 3. The Gaussian kernel is initially broader and then gradually reduced with the iterations [Shen 2002]. As the algorithm progresses, more and more driving voxels are added, increasing the dimensionality of the energy function and thus rendering the matching function less and less smooth. However, local minima are avoided, because the algorithm's starting point is closer to the global minimum, each time driving points are added.

Resolving multiple matches. Consider a voxel that has been chosen to be a driving voxel, at a particular stage of the deformation mechanism. A search neighborhood around that driving voxel is first defined, and the potential candidate matching points are searched, by directly comparing the similarity measures between the attribute vector of that driving voxel and the attribute vectors in all neighboring voxels. If any candidate matches are found, then the whole neighborhood around that driving voxel is tentatively deformed to each matching point, and the new similarity measurement is further calculated by integrating the similarity degrees of all voxels in the neighborhood. This means that the whole anatomy around the driving voxel determines the potential matches. If multiple matches are found this way, then the algorithm does not choose one over the other, but rather it reduces the influence of that particular driving voxel on the deformation. The driving voxel itself deforms in a way that compromises among all candidate matches, until the matching ambiguity is reduced at a later stage of the algorithm. According to this strategy, a sulcus that displays variable topology across individuals will not influence the deformation process significantly, but will primarily follow the deformation of other driving voxels that are less variable.

Consistency of the deformation field. A common concern with morphometric analyses based on shape transformations is the dependency of the transformation on the particular template used each time. At the most fundamental level, a transformation found by an algorithm when transforming brain A to brain B

should be exactly the inverse of what the algorithm finds when it transforms brain B to brain A . In general, this is not the case. In developing HAMMER, we have followed the work of [Christensen, 1999], who introduced the concept of *consistent transformations*, i.e. transformations that are found by enforcing the consistency of the forward and the inverse transformations. This is accomplished by solving a matching function that is symmetric in terms of the forward and the inverse transformation, both of which are found simultaneously and are forced to be consistent with each other. For example, when the displacement of a driving voxel is found, the displacements of all candidate matching points are found simultaneously, as they correspond to the inverse transformation. If the match of point P is determined to be point Q , but the match of point Q is not found to be point P , an average is formed that balances between these two terms and enforces consistency of the transformation on the driving points.

Smoothness of deformation field. In order to ensure that the Jacobian of the deformation field is positive, the deformation field should be smooth enough. Accordingly, we employ several strategies for making our deformation field smooth. Specifically, according to our sub-volume deformation strategy, the neighboring voxels deform smoothly and together with the driving voxel, as determined by a Gaussian kernel. Moreover, during each iteration of the algorithm, two smoothing techniques are used to refine the deformation field. The first technique uses global/local affine transformations, which are estimated from the displacements of the driving voxels by using a linear least square estimation procedure, to constrain the global/local consistency of the displacement fields. This is particularly important in the initial deformation stages, at which time the driving voxels are sparse in the image space and thereby easier to be affected by noise. The second smoothing technique is the use of the Laplacian smoothness constraint. All of these smoothing techniques make the deformation field well behaved, in that they are smooth and have smooth first derivatives [Shen 2002].

III. Results

In this section we describe a series of experiments that we performed to validate HAMMER, and we demonstrate how our approach could be used to quantify local patterns of atrophy, in a sample group of elderly subjects.

Experiment 1. In our first experiment we selected 40 male subjects from the Baltimore Longitudinal Study of Aging [Resnick, 2000], in which we have previously segmented the ventricles by manually outlining the part of CSF that belongs to the ventricles. (CSF was determined via a voxel classification

technique [Yan, 1995], validated in [Goldszal 1998].) In order to evaluate our approach, we used HAMMER for automatic labeling of the ventricles in the same subjects. Automatic labeling was obtained by labeling the template, and by transferring the labels after deformable registration with the subjects. For better spatial specificity in our error measurements, we separated the left from the right ventricles and performed the measurements individually.

The results are shown in Fig. 4 for volumetric measurements. From these results, we can see an almost perfect agreement between automatic and manual segmentation of the ventricles.

Experiment 2. In our second experiment, we randomly selected 18 subjects from the same aging study, with a wide variability in brain atrophy. Typical cross-sections of the 18 subjects are shown in Fig. 5, which shows the variability in brain atrophy. We then applied HAMMER to spatially normalize the MR images of these 18 subjects to the space of a randomly chosen template image, which also belonged to the same study. The spatial transformations were determined from the segmented versions of the same images; segmentations into gray matter, white matter, and CSF were determined via the method described in [Goldszal, 1998]. We finally formed the average of the 18 spatially normalized MR images, which is shown in Fig. 6, revealing a very good alignment of these 18 subjects, as reflected by the sharpness of the average image. We note that the accurate matching of the individual and the template is the key issue in a subsequent morphometric analysis. The sharpness of the average implies that morphometric measurements of the resulting tissue density maps will be able to resolve very small structures and localized effects.

Experiment 3. In our third experiment we synthesized brain deformations, and then tested how well HAMMER was able to retrieve these deformations. A key issue in generating synthesized deformations is to use a realistic transformation. For this reason, we used real brains in order to determine the synthesized deformations. In particular, we selected one brain image to be used as template, and five other brain images to be used to generate five synthesized deformations. We extracted the outer cortical surface in these images, as described in [Davatzikos, 2001b]. We then manually outlined a number of sulcal curves on the outer cortical surface, by the method described in [Davatzikos, 2001b]. These curves included the precentral, central, postcentral, superior temporal, superior frontal, inferior frontal sulci, and Sylvian and interhemispheric fissures. The outlined curves were then used as anatomical features that guided a 3D elastic warping of the template to the 5 other brains. In addition to cortical deformations, we imposed ventricular deformation, by introducing different degrees of ventricular contraction. Contraction was

achieved by introducing a uniform strain within the ventricles, which after relaxation of the elastic warping resulted in varying ventricular sizes [Davatzikos, 1997]. Typical cross-sections and 3D renderings of the resulting brains are shown in Fig. 7.

We then applied HAMMER to spatially normalize the five synthesized brains to the template. We note that HAMMER is a completely independent method from the one that was used to synthesize the deformations. The average of the resulting spatially normalized images is shown in Fig. 8, in triplanar display, and reveals a very good alignment. In addition to this visual inspection of the spatial normalization accuracy, we measured the alignment of two structures after spatial normalization: the right precentral gyrus (PCG) and the left superior temporal gyrus (STG). We outlined these regions in the template, and we deformed their labels according to the synthesized deformations. We then followed the HAMMER-based spatial normalization of these two regions, and measured their degree of overlap among the 5 spatially normalized images in the template space. The result is shown in Fig. 9, in which black corresponds to 100% overlap, blue to nearly 100% overlap, and red to 0% overlap. Slight disagreements of the boundary of these regions are in part due to discretization errors introduced by the binary nature of the labels during the synthesized deformation and the spatial normalization.

Experiment 4. In our fourth experiment, we compared manual versus automatic segmentations of the PCG and the STG in 11 brains. In particular, we randomly selected 11 images from the BLSA study, and we labeled the PCG and STG with two independent raters. We then used HAMMER-based warping of the labeled template image to obtain a fully automated labeling of these two regions. Representative sections showing the labels of the two raters and the one obtained via HAMMER are shown in Fig. 10, with the corresponding volume and overlap errors stated in the bottom of each subfigure. Fig. 11 shows the results for the PCG (left column) and STG (right column). The top row shows plots of the volumetric measurements obtained via the three methods. The middle row shows plots of the percent overlap errors between HAMMER and each rater, and the two raters, and the bottom row shows the respective percent errors in volume measurements. Table 1 summarizes the average (for all 11 subjects) errors for the two structures.

Experiment 5. In our final experiment, we demonstrate the utility of the RAVENS tissue density maps in quantifying local volumetric group differences, using the white matter maps as an illustration. In particular, we generated the RAVENS maps for the 40 subjects described earlier, and we then divided the subjects into two groups, 20 relatively younger elderly (ages 59-68, with average age 61.1) and 20

relatively older elderly (ages 69-84, with average age 74.1). Fig 12a shows the deformations of a typical subject slice to a selected template brain, with the color-coded Jacobian map of the same slice given in Fig 12b. By incorporating several smoothing techniques into HAMMER algorithm, the Jacobian is positive in the whole deformation space, as shown by Fig. 12b. In order to demonstrate how one would measure local volumetric group differences, we subtracted the average RAVENS map of WM of the *older* elderly group from that of the *younger* elderly group, and we display the difference, as color-coded in Fig. 13 and as 3D rendered image in Fig. 14. Notably, we normalized the RAVENS maps by global brain size differences, by first applying an affine transform that removed such differences. Therefore, the color-coded map of Fig. 13 displays the relative volumetric differences between the two groups. For example, the splenium displays relatively higher volumetric difference, a result that is verified by our previous detailed analysis of the corpus callosum [Davatzikos, 1998a]. We note that this application is from a limited sample, but demonstrates the utility of this approach for quantification of local patterns of atrophy. A comprehensive analysis of morphometric effects of aging using the BLSA cohort was presented in [Resnick 2000] using a spatial normalization method of relatively lower spatial specificity, and will be the subject of a follow-up paper, using the approach described herein.

IV. Discussion

This paper presented a methodology for voxel-based morphometry of very high spatial resolution, which is achieved via a very flexible and fully automated spatial normalization approach, referred to as HAMMER. The main novelties of HAMMER are twofold. First, it uses the concept of an attribute vector, i.e. a collection of geometric attributes whose goal is to uniquely characterize every single voxel in a brain image, thereby reducing ambiguity in the matching process. Second, it uses a hierarchical approximation of the similarity function, thereby significantly reducing local minima, which typically represent poor matches. HAMMER is used in conjunction with a mass-preserving mechanism, originally presented in [Goldszal, 1998; Davatzikos, 1998b; Davatzikos, 2001b], which ensures that no volumetric information is lost during the process of spatial normalization, since this process changes an individual's brain morphology in order to conform it to the morphology of a template.

This approach was validated in five experiments. In our first experiment, we compared the segmentation and volumetric measurement of the lateral ventricles in a group of older adult men, some of whom displayed significant brain atrophy. The scatterplots of Fig 4 show an almost perfect agreement between automatic and manual segmentation of the ventricles. Some minor differences were observed only in the

subjects that have extremely large ventricles, revealing a potential limitation of this approach in cases of extreme atrophy.

Our second experiment visually demonstrated the accuracy of the spatial normalization, via the formation of an average image resulting from 18 spatially normalized T1-weighted images of older adults. The sharpness of the average of the warped brain images is often used as a visual display of the accuracy of the warping algorithm, although the sharpness of average images was not found to be a very good indicator of registration accuracy for the evaluations in [Hellier, 2001]. Fig. 6 shows an exquisite sharpness, indicative of the registration accuracy of the underlying normalized images. As expected, particularly good is the alignment of the subcortical structures, which resulted in improved signal to noise characteristics of the average image, compared to the individual images. This can be seen, for example, at the claustrum, the putamen, and the thalamus. It is important to note that there is residual variability in the cortex, which is primarily due to fundamental morphological differences between the template and the individuals. Typical examples are the “one versus two” sulci, or the absence of a sulcus in an individual, which is present in the template. In such cases, HAMMER is designed to relax the matching forces when no good matches are found. That is, a gyrus or a flat part of the cortex will not be forced to match a sulcus, simply because a sulcus happens to be in its vicinity. In order for a match to be enforced, a high similarity of the attribute vectors must be present. Otherwise, the matching will be driven by other features, which display high degree of similarity of their attribute vectors.

Although sharpness of an average image does not directly demonstrate accuracy in morphometric measurements, accuracy of the spatial normalization is a key issue. This is because a voxel-wise comparison and statistical analysis of the resulting tissue density RAVENS maps assumes that a given voxel is compared with its corresponding voxel in all subjects. In principle, this assumption is violated whenever the spatial normalization error is of the order of or higher than the dimensions of half a voxel. In practice, the smoothness of the warping fields somewhat relieves this almost unachievable requirement. Nonetheless, spatial normalization accuracy directly affects the sensitivity of a subsequent voxel-wise statistical analysis of the RAVENS maps.

Our third experiment demonstrated that HAMMER can capture synthesized deformations that were based on five different brains. Using synthesized deformations to evaluate the accuracy of a spatial normalization method needs to be performed with caution, since an unrealistic synthesized deformation may not be recoverable, which is not necessarily a limitation of the spatial normalization algorithm being

evaluated. Conversely, an unrealistic synthesized deformation can be “too easy”, and easily recoverable by the spatial normalization algorithm, without this being indicative of a more general accuracy of the algorithm. We made an effort to address this issue, by using real brain images to synthesize the deformations, thereby generating synthetic brains that reflect real variability. Admittedly, we only used a dozen of cortical features to synthesize these deformations, which are hardly adequate to capture the variability of the cortex. However, it is important to note that the two regions that we measured, namely the PCG and the STG, were surrounded by cortical features used to generate the synthetic deformations. Therefore, the true variability of these particular gyri was captured in the generation of these synthetic deformations.

Our fourth experiment compared the segmentations that were obtained automatically by warping a template brain to each of 11 individual brain images, for the PCG and the STG, with the respective manual labelings of these gyri. We demonstrated that the difference between the algorithm and the two raters was the same or even smaller than the difference between the two raters. This implies that the segmentations of the algorithm were somewhere in-between those of the two raters. Notable in that experiment was the high variability in the measurements. This is mostly due to difficulties in defining the boundaries of these gyri, especially towards the white matter, which are rather arbitrary. In the face of such arbitrary definitions, an additional advantage of HAMMER is that these definitions will be applied with high reproducibility.

Finally, our fifth experiment demonstrated the use of the tissue density RAVENS maps and Hammer algorithm in capturing local volumetric group differences. A more extensive analysis of the morphological effects of aging, using the approach described in this paper, is part of our future work.

One of the apparent limitations of the approach is the fact that the density maps appear to be slightly “noisier” than one should perhaps desire. Although we have no underlying truth to support this claim, one would expect higher spatial smoothness in the group differences measured in Experiment 5. A manifestation of this limitation is the presence of several small and isolated blue regions in Fig. 13, which imply age differences in volumes. Work on better models for the spatial transformation, and in particular for interpolating the deformation in-between driving features, will alleviate this limitation.

The RAVENS maps are suitable only for local volumetric analysis, and not for other types of shape measurements. For example, two structures might have the same volumes, both globally and locally, but

one might have higher curvature than the other. Our main focus in the approach presented herein has been on local volumetrics, since they can be interpreted as local loss or growth of brain tissue. It is unclear what a difference in the curvature of two structures might mean, or if such curvature differences are caused by the loss or growth of adjacent brain tissue, rather than by processes taking place within the structure. However, our approach is directly generalizable, as it can measure the shape transformation that maps the template to the individual, which captures all morphological characteristics of the anatomy being measured. Several groups [Thompson, 2001; Miller, 1997; Joshi, 1997; Bookstein, 1989], including ours [Davatzikos, 1996; Davatzikos, 2001a], have taken similar approaches.

The current attribute vector used by HAMMER has relatively limited discriminating power. For example, sulci are clearly distinguished from gyri, but the precentral sulcus cannot be distinguished from the postcentral sulcus, based purely on the attribute vector that is calculated from a relatively small neighborhood. Current and future work in our laboratory includes the construction of a richer attribute vector with much higher discriminating power. Ideally, if every voxel in the brain has a unique and robust morphological signature, finding the matching transformation will be a trivial issue.

V. Acknowledgments

We would like to thank Dr. Susan Resnick and the BLSA for providing the datasets. This work was supported in part by NIH grant R01 AG14971 and by NIH contract AG-93-07.

References:

- [Ashburner, 1998] J. Ashburner, C. Hutton, R. Frackowiak, I. Johnsrude, C. Price, and K. Friston. "Identifying global anatomical differences: deformation-based morphometry". *Hum. Brain Mapp*, **6**: 348–357, 1998.
- [Ashburner, 2000] J. Ashburner and K.J. Friston. "Voxel-based morphometry: the methods". *NeuroImage*, 11: 805-821, 2000.
- [Bajcsy, 1983] R. Bajcsy and R. Lieberman and M. Reivich. "A Computerized System for the Elastic Matching of Deformed Radiographic Images to Idealized Atlas Images". *Journal of Computer Assisted Tomography*, volume 7, number 4, p.618-625, 1983.
- [Bookstein, 1989] Bookstein FL. "Principal warps: Thin-plate splines and the decomposition of deformations". *IEEE Trans Pattern Anal Mach Intell*, 11:567–585, 1989.
- [Bookstein, 2001] F. L. Bookstein. "Voxel-Based Morphometry Should Not Be Used with Imperfectly Registered Images", *NeuroImage*, 14:1454-1462, 2001.
- [Breijl, 2000] 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*, volume 19, p.973-985, 2000.
- [Briquer, 1997] Briquer, L.L., and Gee, J.C.. "Design of a statistical model for brain shape". *Lect. Notes Comp. Sci.:Proc. Inform. Proc. Med. Imaging*, 1230: 477–482, 1997.
- [Chen 1999] M. Chen, T. Kanade, D. Pomerleau, J. Schneider, "3-D deformable registration of medical images using a statistical atlas", *MICCAI*, pages 621-630, Sept. 1999.
- [Christensen, 1999] G.E. Christensen. "Consistent Linear-Elastic Transformations for Image Matching". *Information Processing in Medical Imaging*, LCNS 1613, Springer-Verlag, pp. 224-237, 1999.
- [Christensen, 2001] Christensen GE, Johnson HJ. "Consistent image registration". *IEEE Transactions On Medical Imaging*, 20 (7): 568-582 JUL 2001.
- [Chui, 2001] H. Chui, L. Win, R. Schultz, J. Duncan, and A. Rangarajan. "A Unified Feature Registration Method for Brain Mapping". *Information Processing in Medical Imaging*, p.300-314, Davis, CA, USA, June 18-22, 2001.
- [Collins, 1999] Collins DL, Zijdenbos AP, Baare WFC, Evans AC. "ANIMAL+INSECT: Improved cortical structure segmentation". *Information Processing In Medical Imaging, Proceedings Lecture Notes In Computer Science*, 1613: 210-223 1999.
- [Davatzikos, 1996] Davatzikos, C., Vaillant, M., Resnick, S., Prince, J. L., Letovsky, S. I., and Bryan, R. N.. "A computerized method for morphological analysis of the corpus callosum". *J. Comput. Assisted Tomogr.* 20: 88–97, 1996.
- [Davatzikos, 1997] C. Davatzikos. "Spatial Transformation and Registration of Brain Images Using Elastically Deformable Models", *Comp. Vis. and Image Understanding*, 66(2):207-222, May 1997.
- [Davatzikos, 1998a] Davatzikos, C., and Resnick, S. M.. "Sex differences in interhemispheric connectivity: Correlations with cognition in women but not in men". *Cereb. Cortex* **8**: 635–640, 1998.

- [Davatzikos, 1998b] Davatzikos, C.. “Mapping of image data to stereotaxic spaces: Applications to brain mapping”. *Hum. Brain Mapp.* **6**: 334–338, 1998.
- [Davatzikos, 2001a] C. Davatzikos. “Measuring biological shape using geometry-based shape transformations”. *Image and Vision Computing*, 19:63-74, 2001.
- [Davatzikos, 2001b] C. Davatzikos, A. Genc, D. Xu, S.M. Resnick, “Voxel-based Morphometry Using the RAVENS Maps: Methods and Validation Using Simulation of Longitudinal Atrophy”, *Neuroimage*, 14:1361-1369, December 1, 2001.
- [Evans, 1991] 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.
- [Freeborough, 1998] P. A. Freeborough and N. C. Fox. “Modeling Brain Deformations in Alzheimer Disease by Fluid Registration of Serial MR Images”, *Journal of Computer Assisted Tomography*, 22:838-843, 1998.
- [Gee, 1993] 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.
- [Gee, 1994] 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*, vol. 2359, pp.191-204, 1994.
- [Goldszal, 1998] 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.
- [Hellier, 1999] Hellier P, Barillot C, Memin E, Perez P. “Medical Image Registration With Robust Multigrad Techniques”. *MICCAI'99, Proceedings Lecture Notes In Computer Science*, 1679: 680-687 1999.
- [Hellier, 2001] P. Hellier, C. Barillot, I. Corouge, B. Gibaud, G. Le Goualher, D. L. Collins, A. C. Evans, G. Malandain and N. Ayache. “Retrospective Evaluation of Inter-subject Brain Registration”, in *Proc. MIUA 2001*, pp. 258-265, Edited by W. J. Niessen and M. A. Viergever. Springer, Utrecht, The Netherlands.
- [Joshi, 1995] 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*, volume 2573, p.278-289, July 1995.
- [Joshi, 1997] S.C. Joshi, M.I. Miller, U. Grenander. “On the geometry and shape of brain submanifolds”. *Intl. J. of Pattern Recognition and Artificial Intelligence*, 11: 1317-1343, 1997.
- [Machado, 1998] A. M. C. Machado and J. C. Gee. “Atlas warping for brain morphometry”, in *Proc. SPIE Medical Imaging 1998: Image Processing*, K. M. Hanson, ed., Bellingham, WA: SPIE.
- [Miller, 1993] M.I. Miller, G.E. Christensen, Y. Amit, U. Grenander. “Mathematical textbook of deformable neuroanatomies”. *Proc Nat Acad Sci USA*, 90:11944–11948, 1993.
- [Miller, 1997] M. Miller, A. Banerjee, G. Christensen, S. Joshi, N. Khaneja, U. Grenander, L. Matejic. “Statistical methods in computational anatomy”. *Statistical Methods in Medical Research*, 6: 267-299, 1997.

- [Pizer, 1999] S.M. Pizer, D.S. Fritsch, P. A. Yushkevich, V. E. Johnson, and E. L. Chaney, "Segmentation, Registration, and Measurement of Shape Variation via Image Object Shape", *IEEE Trans. on Medical Imaging*, 18(10): 851-865, October 1999.
- [Resnick, 2000] S.M. Resnick, A.F. Goldszal, C. Davatzikos, S. Golski, M.A. Kraut, E.J. Metter, R.N. Bryan, and A.B. Zonderman, "One-year age changes in MRI brain volumes in older adults", *Cerebral Cortex*, 10:464-472, 2000.
- [Rohr, 1999] K. Rohr. "Image registration based on thin plate splines and local estimates of anisotropic landmark localization uncertainties". *Lect. Notes in Comp. Sci.:MICCAI'98*, Volume 1496, p.1174-1183, 1999.
- [Rueckert, 1999] Rueckert D, Sonoda LI, Hayes C, Hill DLG, Leach MO, Hawkes DJ. "Nonrigid registration using free-form deformations: Application to breast MR images", *IEEE Trans. on Medical Imaging*, 18 (8): 712-721 AUG 1999.
- [Shen, 2001] D. Shen, C. Davatzikos, "HAMMER: Hierarchical Attribute Matching Mechanism for Elastic Registration", *IEEE Workshop on Mathematical Methods in Biomedical Image*, Kauai, Hawaii, USA, December 9-10, 2001.
- [Shen, 2002] D. Shen, C. Davatzikos, "HAMMER: Hierarchical Attribute Matching Mechanism for Elastic Registration", *IEEE Trans. on Medical Imaging*, 21(8), August 2002.
- [Stefansic, 2000] Stefansic J.D., Herline A.J., Shyr Y., Chapman W.C., Fitzpatrick J.M., Dawant B.M., and Galloway R.L. "Registration of physical space to laparoscopic image for use in minimally invasive hepatic surgery", *IEEE Transactions on Medical Imaging*, 19(10), pp. 1012-1023, 2000.
- [Subsol, 1998] G. Subsol, J. Ph. Thirion, N. Ayache. "A general scheme for automatically building 3D morphometric anatomical atlases: application to a skull atlas". *Medical Image Analysis*, 2(1):37-60, 1998
- [Thévenaz, 2000] P. Thévenaz, M. Unser, "Optimization of Mutual Information for Multiresolution Image Registration," *IEEE Transactions on Image Processing*, vol. 9, no. 12, pp. 2083-2099, December 2000.
- [Thirion, 1992] J.P. Thirion, O. Monga, S. Benayoun, A. Gueziec and N. Ayache. "Automatic registration of 3-D images using surface curvature". *SPIE Proc., Mathematical Methods in Medical Imaging*, volume 1768, p.206-216, 1992.
- [Thirion, 1996] J.P. Thirion. "Non-rigid matching using deamons". *Proc. IEEE Conf. Comp. Vis. and Patt. Recog.*, 1996.
- [Thompson, 1997] Thompson, P. M., MacDonald, D., Mega, M. S., Holmes, C. J., Evans, A. C., and Toga, A. W.. "Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces". *J. Comput. Assisted Tomogr.* 21: 567-581, 1997.
- [Thompson, 2000] Thompson, P. M., Woods, R. P., Mega, M. S., and Toga, A. W.. "Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain". *Hum. Brain Mapp.* 9:81-92, 2000.
- [Thompson, 2001] Thompson, P.M., Mega, M.S., Vidal, C., Rapoport, J.L., Toga, A.W. (2001). "Detecting Disease-Specific Patterns of Brain Structure using Cortical Pattern Matching and a Population-Based Probabilistic Brain Atlas", *Proc. 17th International Conference on Information Processing in Medical Imaging (IPMI2001)*, Davis, CA, USA, June 18-22, 2001, pp. 488-501.

[Vemuri, 2001] B. C. Vemuri, J. Liu, and J.L. Marroquin. "Robust Multimodal Image Registration Using Local Frequency Representations". *Information Processing in Medical Imaging*, p.176-182, Davis, CA, USA, June 18-22, 2001.

[Wang 2000] Y. Wang and L. H. Staib, "Physical model-based non-rigid registration incorporating statistical shape information", *Medical Image Analysis*, vol. 4, no. 1, pages 7-20, 2000.

[Wright, 1995] I. C. Wright, P. K. McGuire, J.-B. Poline, J. M. Travers, R. M. Murray, C. D. Frith, R. S. J. Frackowiak and K. J. Friston. "A Voxel-Based Method for the Statistical Analysis of Gray and White Matter Density Applied to Schizophrenia", *NeuroImage*, 2:244-252, 1995.

[Yan, 1995] M.X.H. Yan and J.S. Karp. "An adaptive Bayesian approach to three-dimensional MR image segmentation". *Proc. of the Conf. on Inform. Proc. in Med. Imag.*, 1995, pp. 201-213.

Table 1. Percent volume and overlap errors between automatic (HAMMER) and rater-defined labels of the superior temporal gyrus (STG) and the precentral gyrus (PCG) in 11 elderly subjects.

	STG	PCG
<i>Percent Overlap Errors</i>		
HAMMER/Rater1	23.8%	26.1%
HAMMER/Rater2	24.9%	23.8%
Two Raters	23.6%	21.6%
<i>Percent Volume Errors</i>		
HAMMER/Rater1	13.7%	20.9%
HAMMER/Rater2	16.4%	14.6%
Two Raters	24.3%	18.9%

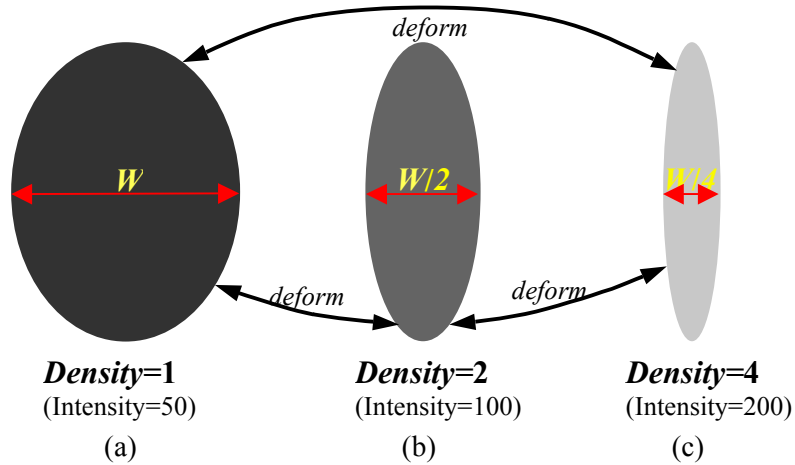


Fig. 1. Schematic demonstration of the mass-preserving deformation principle. As pointed by the black solid arrow, squashing of the object in (a) to the half width doubles the material density of the object, as shown in (b). Similarly, stretching of the object in (b) to a configuration with double width, halves the material density of the object, as shown in (a). An image intensity value of 50 is used to represent a density value equal to 1 (original object), and image intensity 100 represents double that density (compressed object). Additional compression will further increase density, as in (c).

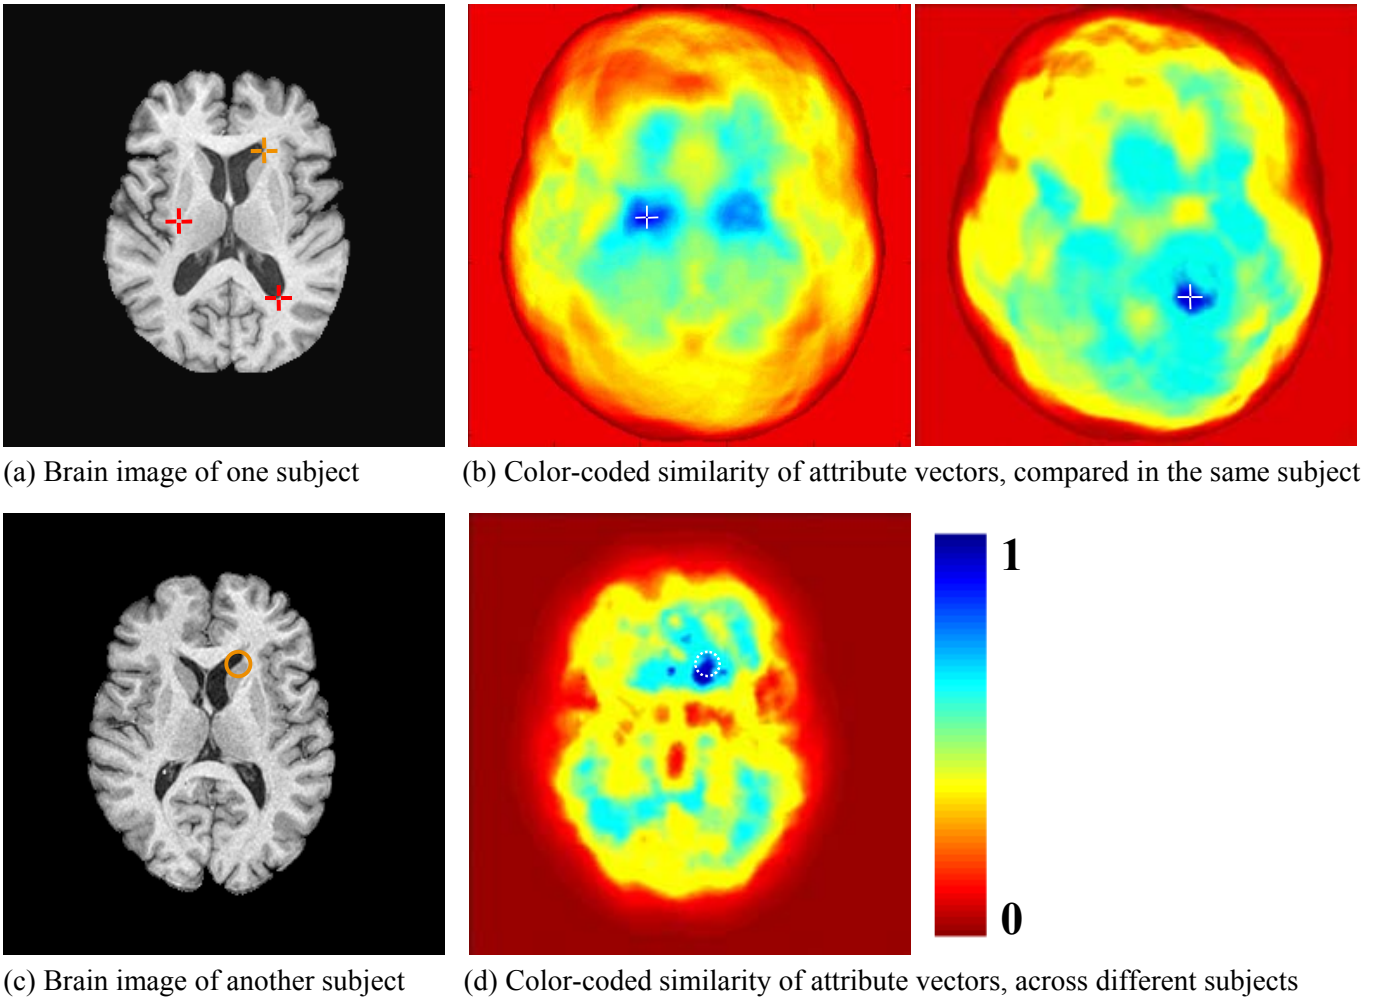



Fig. 2 Demonstration of the geometric moment invariants (GMIs) in discriminating local structures. Fig (a) shows a brain image of one subject. The GMIs of the two voxels indicated by the red crosses are respectively compared with the GMIs of all other voxels in the same subject. Two resulting similarity maps of the GMIs are shown color-coded in (b), with high similarity as blue. The color bar is shown in (d). The two white crosses in (b) correspond to the two red crosses in (a). For the voxel indicated by the orange cross in (a), its GMIs are compared with the GMIs of all voxels in another subject (c). The resulting colored-coded image of the similarity measurements is shown in (d), with the most similar attribute vector in the position that is enclosed by the white dashed circle. The orange circle in (c) has the same coordinates as the white dashed circle in (d), which means that the position enclosed by the orange circle is the corresponding point to the voxel indicated by the orange cross in (a). For all these results, the image size is 256x256; the GMIs are calculated in 81x81 neighborhood for (b) and 51x51 neighborhood for (d). This figure demonstrates that the attribute vector is able to distinguish one voxel from other voxels that could have similar intensity or edge type.



(a) Original template

(b) Template deformed due to a driving voxel

Fig. 3. Demonstration of the Gaussian deformation mechanism used in HAMMER. Figure (a) shows a regular grid and a yellow driving voxel, overlaid on the selected part of the template. The yellow driving voxel deforms to the blue voxel along the direction of the pink arrow. Then the template and the grid are deformed as shown in (b). Yellow and blue voxels, and the pink arrows in (b) are the same as those in (a).

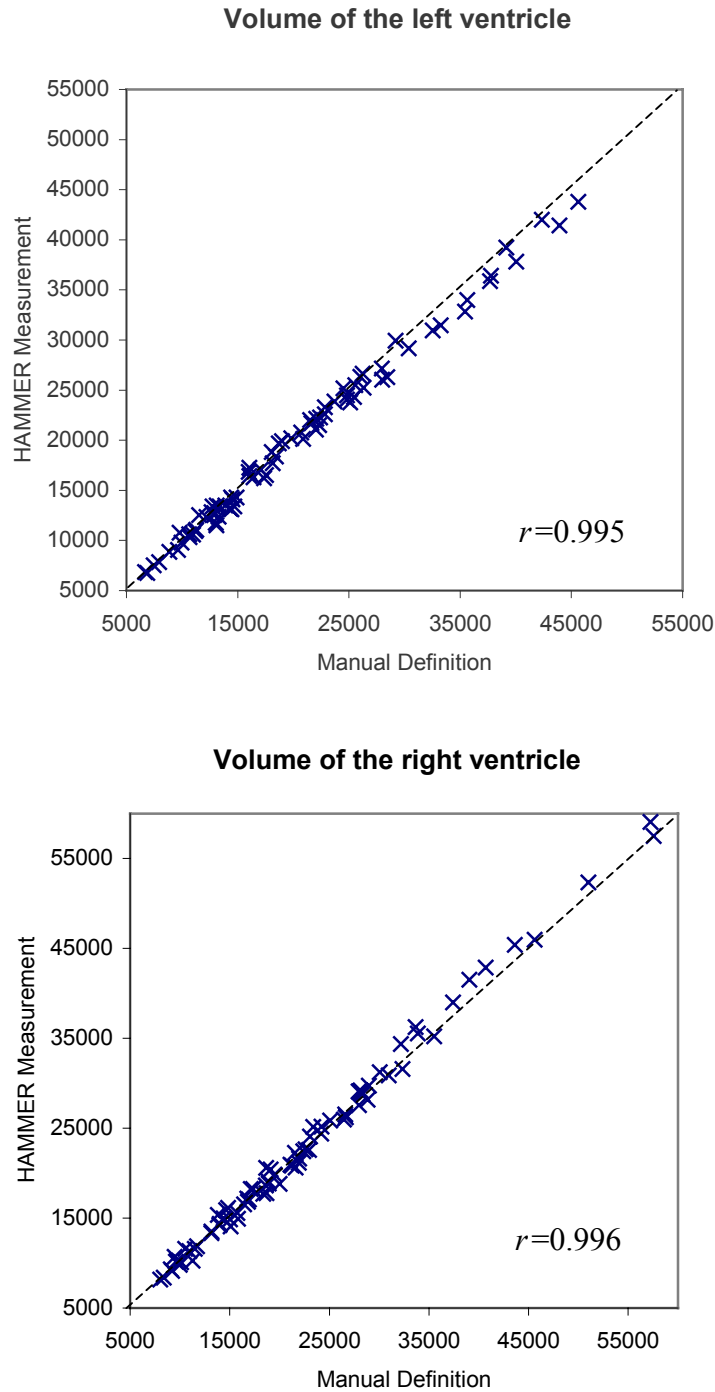


Fig. 4. Comparison on the performance of automatic and manual segmentation of the lateral ventricles in a group of older adults, some of which displayed significant brain atrophy (see Fig. 5, for a representative sample). In each figure, the horizontal axis represents the results of the manual segmentation, and the vertical axis represents the results of the HAMMER-based automatic segmentation. All numbers are in mm^3 . Parameter r is the correlation coefficient between HAMMER measurement and manual definition. This figure shows an almost perfect agreement between automatic and manual segmentation of the ventricles.





Fig. 5. Typical cross-sections of 18 subjects, which show the variability in brain atrophy. These 18 subjects are used to construct the average brain image, after elastic registration.



(a) Average brain by HAMMER

(b) Template brain

Fig. 6. The average brain of the 18 spatially normalized T1-weighted images of elderly subjects. The sharpness of the average brain indicates the very high registration accuracy of the underlying normalized images by HAMMER. Notable is the increase in signal to noise ratio obtained after averaging, which even helps reveal thalamic nuclei.




Fig. 7. Typical cross-sections and 3D renderings of the synthesized brains, used to validate the registration performance of HAMMER. (a) A selected subject; (b-f) five synthesized brains.




Fig. 8. The average brain of five spatially normalized images, resulting from five synthesized brains in Fig 7. This figure is used for the visual inspection of the spatial normalization accuracy by HAMMER.

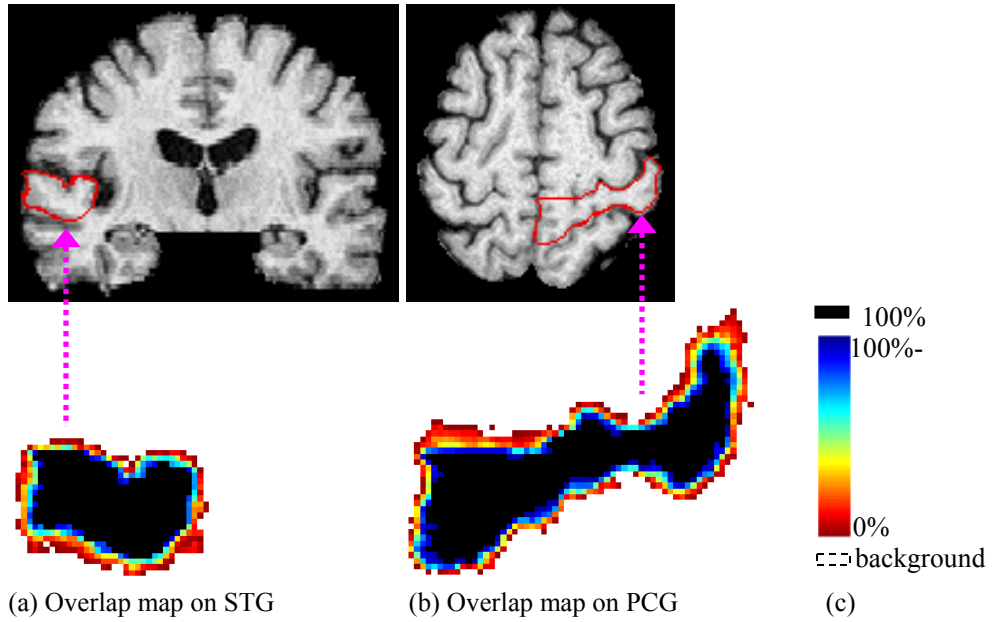
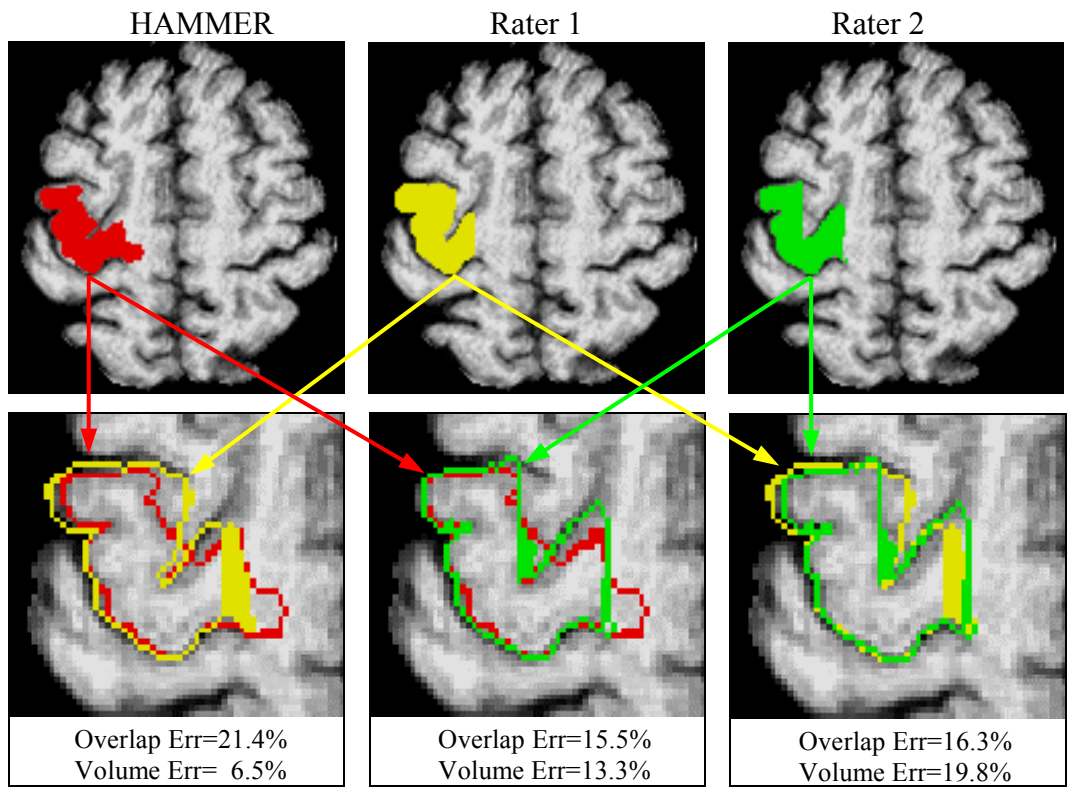
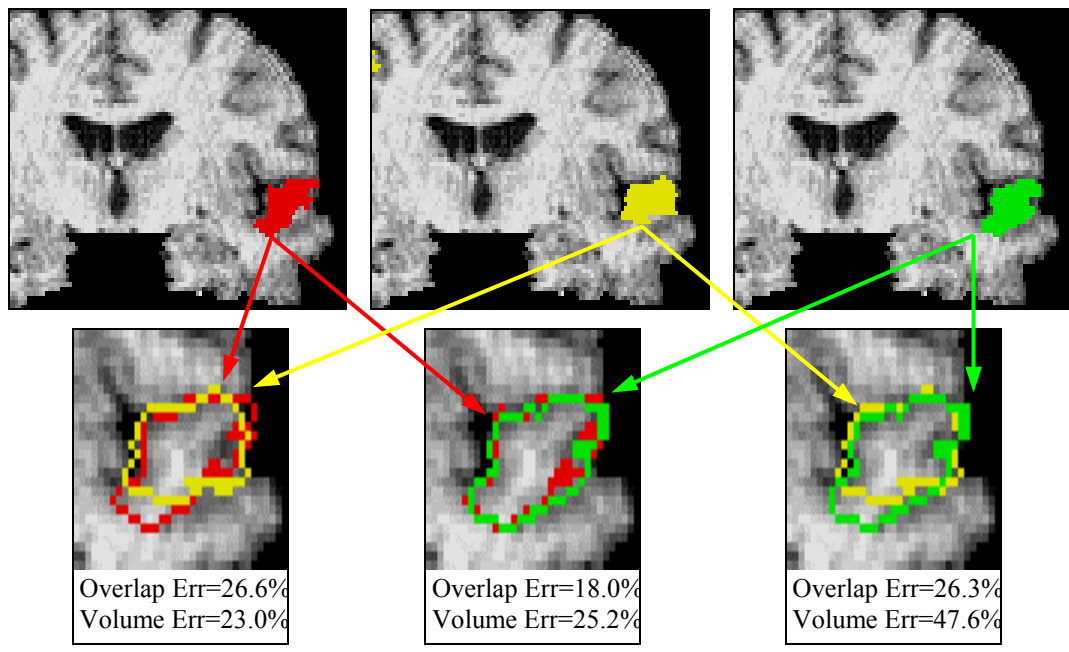


Fig. 9. The alignment of two structures after spatial normalization: the right precentral gyrus (PCG) and the left superior temporal gyrus (STG). Figure (a) gives the overlap map of STG in the template space, and figure (b) gives the overlap map on PCG. Black in the bottom of (a) and (b) corresponds to 100% overlap, blue to nearly 100% overlap, and red to 0% overlap. The detailed color-coding bar, with the overlap percentages, is shown in (c).



(a) PCG



(b) STG

Fig. 10. Representative sections showing the labels of the two raters and the one obtained via HAMMER. The left column shows the automatic labeling results by HAMMER. The middle and the right columns respectively show the results by two raters. Figure (a) shows the labels on PCG, and figure (b) shows the labels on STG. The percent volume and overlap errors are stated in the bottom of subfigures (a) and (b) as examples for the reader to better interpret the results of Table 1.

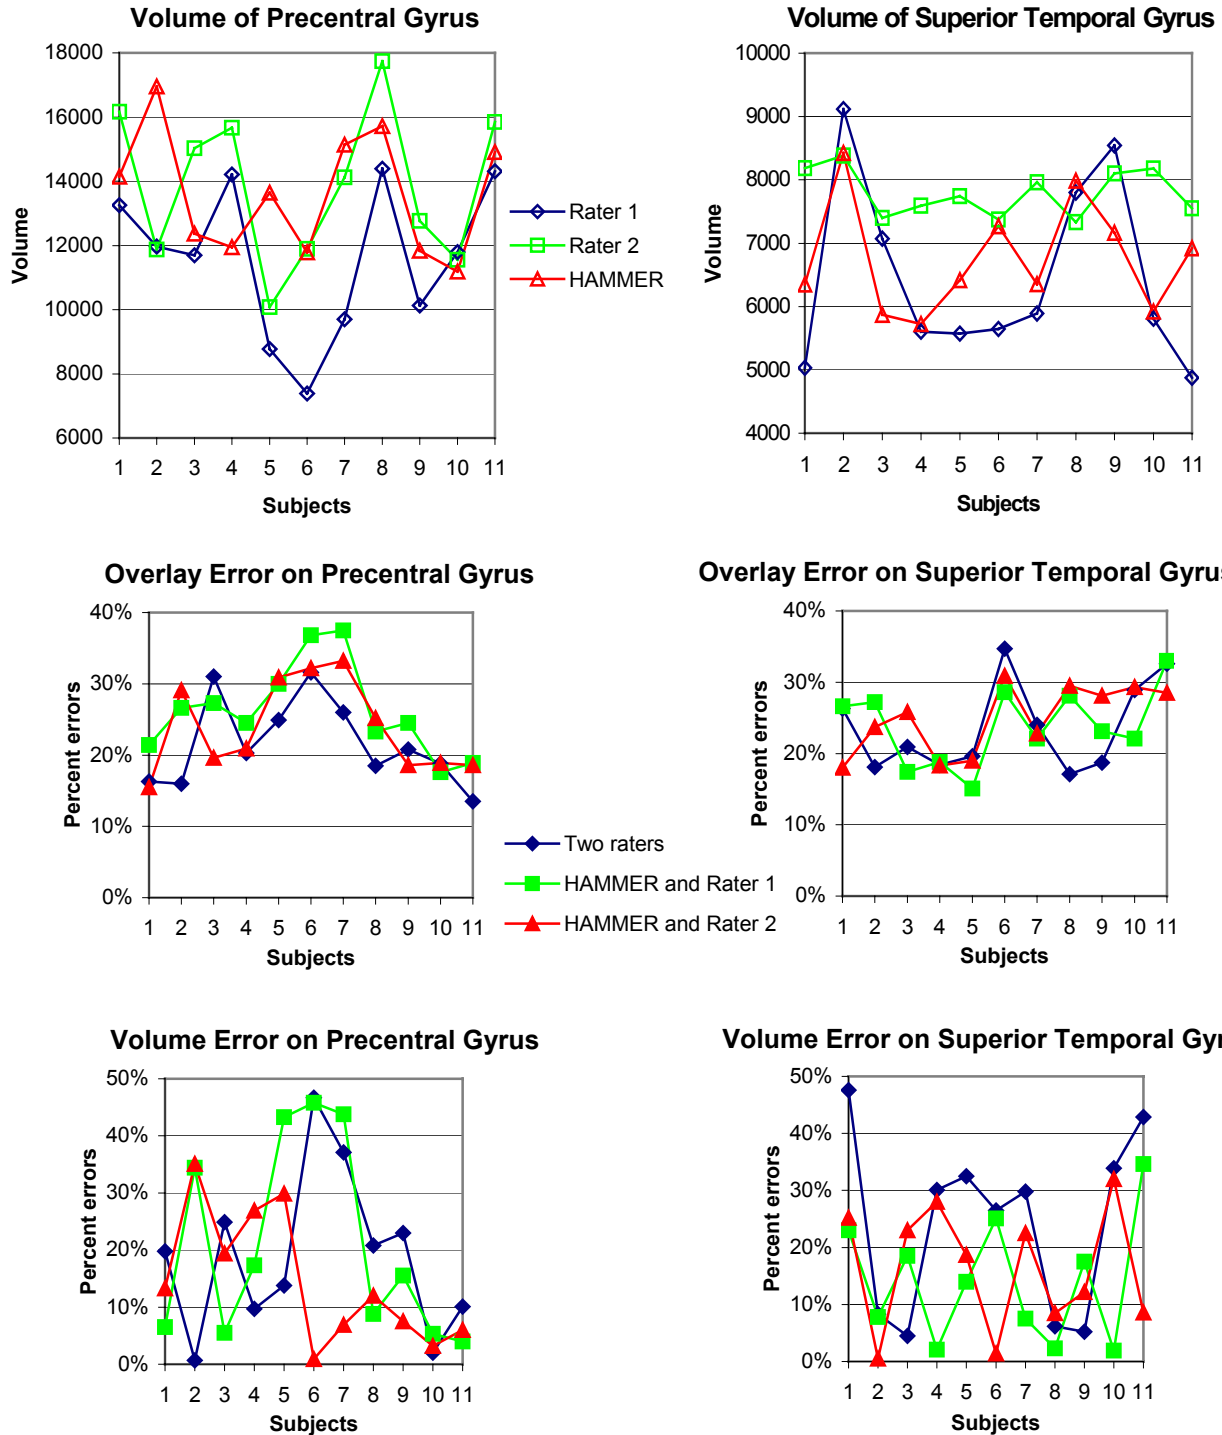


Fig. 11. Comparison of manual versus automatic segmentations of the PCG and the STG in 11 brains. The top row shows plots of the volumetric measurements obtained via the three methods, HAMMER and two independent raters. The unit for the volumes is mm^3 . The middle row shows the percent overlap errors between HAMMER and each rater, and the two raters, and the bottom row shows the respective percent volume errors. The difference between the algorithm and the two raters was the same or even smaller than the difference between the two raters.




Fig. 12. (a) Deformation of a Cartesian grid for a typical slice of a BLSA subject. (b) Color-coded map of the determinant of the Jacobian of the deformation in (a). The white contours in (b) are the white matter boundaries of the brain on the left. The Jacobian here reflects differences between the subject and the template. Note that the Jacobian is always positive.




Fig. 13. Demonstration of the way in which the RAVENS tissue density maps could be used to quantify local volumetric group differences. This is the color-coded difference between the average RAVENS maps of WM of the two groups, (*younger elderly group - older elderly group*). The unit for the number in the color bars is mm^3 per voxel; the size of each voxel is $0.9375 \times 0.9375 \times 1.5 \text{ mm}^3$. For the selected planar region of corpus callosum, the total volumetric change was 52.5 mm^3 , which corresponds to 56 mm^2 area difference. This shows how ROI-based measurements can be obtained by integrating the RAVENS maps within an ROI defined on the template.




Fig. 14. 3D renderings of local WM volumetric group differences, (*younger elderly group - older elderly group*). Yellow colors indicate the regions of large volume loss. The cross-sectional version of this result is displayed in Fig. 13.