

Quantitative Study of Brain Anatomy

Mei Chen, Takeo Kanade, Henry A. Rowley, Dean Pomerleau

CMU-RI-TR-98-05

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

March, 1998

© 1998 Carnegie Mellon University

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. government.

Abstract

We introduce a system that automatically segments and classifies features in 3-D images. The system's accuracy is comparable to manual segmentation. It takes 12 minutes to segment and classify 144 brain structures in 256x256x124 voxel image, while similar work by human took 8 months.

The process starts with an atlas, a hand segmented and classified MRI of a normal brain. Given a subject's data, the atlas is warped in 3-D using a hierarchical deformable matching algorithm until it closely matches the subject, i.e. the atlas is customized for the subject. The customized atlas contains the segmentation and classification of the subject's anatomical structures.

The system has processed MRI of 105 subjects, and for 97 of them produced segmentation qualitatively comparable to manual segmentation. We performed quantitative evaluation of the precision for one structure. Of 18 subjects for which classification correctness was examined voxel by voxel, an error rate of less than 20% was achieved for 12 subjects, and less than 10% for 5 subjects.

The efficiency, accuracy, and consistency of the system's performance allow for detailed quantitative studies of brain structures. Initial results have been obtained for finding the normal range of variation in the size and symmetry properties of anatomical structures, and for detecting abnormalities.

Table of Contents

1. Motivation

2. Automated Segmentation and Labelling System

2.1 Similarity Transformation

2.2 Smooth Deformation

2.3 Fine-tuning Deformation

3. Performance Evaluation

4. Statistical Study of Brain Structures

4.1 The Range of Normal Variations

4.2 Symmetry within the Human Brain

5. Application of the Statistics

6. Related Work

7. Discussion and Future Directions

Acknowledgments

References

Quantitative Study of Brain Anatomy

Mei Chen
meichen@cs.cmu.edu

Takeo Kanade
tk@cs.cmu.edu

Henry A. Rowley
har@cs.cmu.edu

Dean Pomerleau
pomerlea@cs.cmu.edu

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

We introduce a system that automatically segments and classifies features in brain MRIs. It takes 22 minutes to segment 144 structures in a $256 \times 256 \times 124$ voxel image on an SGI computer with three 194 MHz R10K processors. The accuracy is comparable to manual segmentation, which can take an expert at least 8 months.

The process starts with an atlas, a hand segmented and classified MRI of a normal brain. Given a subject's data, the atlas is warped in 3-D using a hierarchical deformable matching algorithm until it closely matches the subject, i.e. the atlas is customized for the subject. The customized atlas contains the segmentation and classification of the subject's anatomical structures.

Qualitative and quantitative evaluations of the system's performance show promise for applications in the quantitative study of brain anatomy. We have obtained initial results for finding the normal range of variation in the size and symmetry properties of anatomical structures, and for detecting pathologies.

1. Motivation

In many medical applications, such as the study of Alzheimer's disease and schizophrenia, it is important to examine the relevant anatomical structures quantitatively. Developments in medical imaging (e.g. magnetic resonance imaging, MRI) have made it possible to visualize internal anatomical structures. However, due to the lack of an automated segmentation and labelling system with sufficient accuracy and efficiency, most quantitative assessment of individual structures has been done manually.

One example is a study of physically abused children. To examine the volumes of children's brain structures, the researchers have hired operators who laboriously label MRI volumes, slice by slice, using a computer interface. For each 2-D slice, the operator circles an area that contains the structure of interest. The interface displays the intensity histogram of that area. The operator then picks a threshold based on the histogram. Segmentation by that threshold gives the 2-D area of the structure in that slice. The 3-D volume of the structure is computed by summing up the structure's 2-D area in all the slices. To examine all the brain structures in

this manner will take at least a year, which makes study of many subjects prohibitive. Currently, doctors compromise by studying a limited number of structures over a manageable number of subjects. For the study to have sufficient statistical significance, however, a database for more subjects is mandatory.

In addition to being time-consuming, such manual segmentation suffers from two deficiencies. One is that there exist density variations within each anatomical structure, and single thresholds cannot segment item accurately. The other is that the segmentation is based on human judgement, and is prone to be subjective, inconsistent, and non-repeatable, which reduces the credibility of conclusions drawn from them.

We will introduce a system that can automatically segment and classify 144 important brain structures in a $256 \times 256 \times 124$ voxel image in 22 minutes, with an accuracy comparable to manual segmentation, as shown in Figure 1. The left image is one slice of a subject's MRI, the right is the segmentation generated by the system. Different color regions represent different anatomical structures. Once the segmentation is available, many types of quantitative analysis can be conducted. As shown in the figure, the user can choose any structure, and the system instantly overlays the segmentation on the MRI data, and displays the name, volume, and average density of the structure.

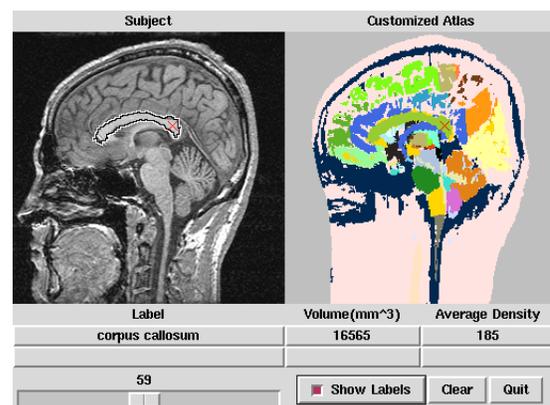


Figure 1: A subject's data (left), and its automatic segmentation (right). When the user selects a structure, the segmentation is overlaid on the raw data. The structure's name, volume, and average density are given at the bottom.

The following sections will discuss how the system works, its performance, and examples of its applications. Please refer to [22] for details on the algorithm and implementation issues.

2. Automatic Segmentation and Classification

The system achieves automatic segmentation and classification by taking advantage of an *atlas*, i.e. an MRI with known segmentation and classification. When imaged using the same modality, the corresponding anatomical structures in the atlas and a subject resemble each other in appearance. This system applies a hierarchy of deformable models to the atlas, warping it in 3-D space to match the subject's data. Once the atlas fits the subject's data, the segmentation and classification in the atlas directly applies to the subject, i.e. the atlas is *customized* for the subject.

The system works on T1 weighted magnetic resonance images (MRI). An MRI volume is a series of parallel cross-sections along one of three principal axes: coronal, axial, or sagittal. Example slices of a brain MRI taken along each axis are shown in Figure 2.

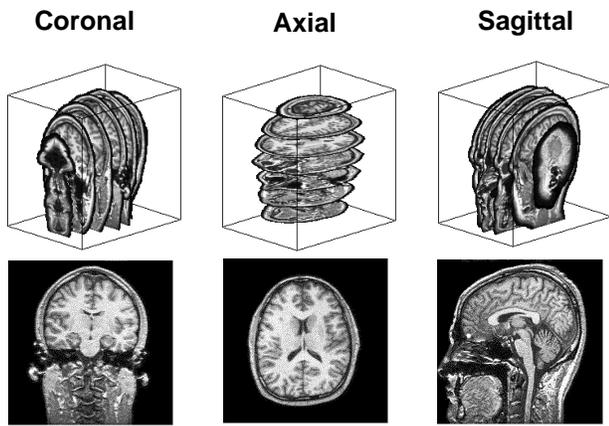


Figure 2: Example slices of brain MRI along each principal axis.

Figure 3 shows one slice of the human brain atlas we use. It is a T1 weighted MRI volume of a normal brain. It contains 123 slices, and each slice is a 256x256 pixel matrix. An expert manually segmented 144 structures, and gave each structure a unique label. The labels of the structures were color-coded to illustrate the segmentation.

The atlas may differ from a particular subject's data in two ways. One is the innate difference between people. This causes variations in the shape, size, density, and location of anatomical structures, i.e. variations that are local and intrinsic. The other is the lack of standards in the data acquisition process. The axis along which the images are taken is generally at an angle to the principal axis, and the spacings between consecutive cross-sections may be non-uniform. An MRI can focus on a sub-section of the brain if so desired, and

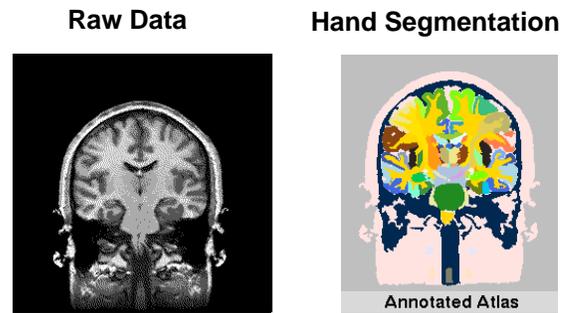


Figure 3: The brain atlas: a T1 weighted MRI volume of a normal brain (left), and the color-coded map of an expert's segmentation and classification of anatomical structures (right).

there can be inhomogeneities in image intensities. These factors cause variations in the orientation, scale, resolution, and intensity consistency between the atlas and subject, i.e. variations that are global and extrinsic.

These two types of variations have different natures, so the system addresses them separately, using a *hierarchical deformable matching algorithm*. It first applies a *similarity transformation* to the atlas, to address the global and extrinsic variations. It then applies a *smooth deformation* to the atlas, to roughly adjust for the local and intrinsic variations. Finally, it uses a *fine-tuning deformation*, which shifts each voxel to fit the atlas more precisely to the subject.

The following subsections will discuss the algorithm and its performance, using the example of a subject with pathology, shown in Figure 4.

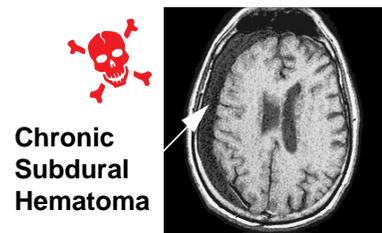


Figure 4: A subject with pathology

2.1. Similarity Transformation

The similarity transformation addresses the global and extrinsic variations between the atlas and the subject by applying a 3-D rotation, uniform scaling, and 3-D translation to the atlas volume. Figure 5 defines the coordinate systems of the volumes. The origins of the coordinate systems in the atlas and the subject's data are placed at their centroids (center of mass). The Z axes coincide with the axes along which the images were scanned. The atlas is first rotated about its centroid, then uniformly scaled about its origin and translated to align with the subject. The similarity transformation T has 7 degrees of freedom.

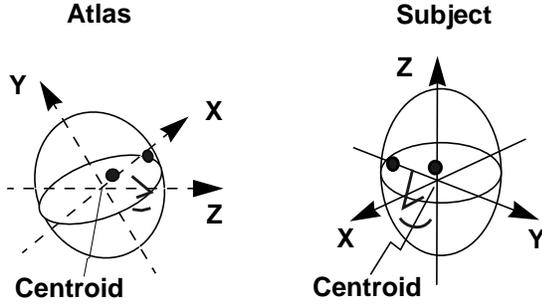


Figure 5: The coordinate systems used in the similarity transformation.

Because the atlas and the subject are inherently different, there is no global transformation that exactly matches them. The best transformation only minimizes the differences. We define the quality of a match in terms of the sum of squared differences (SSD) between the intensities of corresponding voxels in the two volumes. $I_{Subject}(x, y, z)$ is the intensity of voxel $[x, y, z]_{Subject}$ in the subject volume, and $I_{Atlas}(T(x, y, z))$ is the intensity of voxel $[T(x, y, z)]_{Atlas}$ in the atlas volume. $T(x, y, z)$ is (x, y, z) after applying the similarity transformation T. The SSD is a function of this transformation T:

$$ssd = \sum_{(x, y, z)} \left(I_{Subject}(x, y, z) - I_{Atlas}(T(x, y, z)) \right)^2$$

We apply the Levenberg-Marquardt non-linear optimization algorithm to iteratively adjust T to reduce the SSD. The iteration continues until the changes in T are below a user defined criteria. The criteria we use requires that the change in 3-D rotation be less than 2 degrees, that the change in 3-D scaling be below 1%, and that the change in 3-D translation be smaller than 1 pixel.

The system employs multi-resolution processing and stochastic sampling for efficiency and to help prevent the optimization from becoming trapped in local minima [18]. Currently it uses 3 resolution levels.

The transformed coordinates, $T(x, y, z)$, may not be integral. In this case tri-linear interpolation is used to determine the voxel's intensity from its eight bounding neighbors. If the transformed coordinates fall outside of the volume, that voxel is ignored.

Figure 6 shows the result of matching the atlas to the subject using the similarity transformation. The atlas is rotated, scaled, and translated to grossly match with the subject volume. Outlines of the lateral ventricles in the resampled atlas are directly overlaid on the subject data. They only roughly align with their counterparts in the subject, and considerable discrepancies remain between their shapes, sizes, and locations.

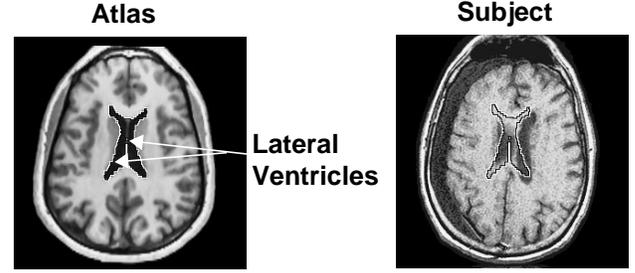


Figure 6: Corresponding slices of the atlas and the subject after the similarity transformation. The atlas is grossly aligned with the subject. The outlines of the lateral ventricles in the atlas are overlaid on the subject's data.

2.2. Smooth Deformation

Smooth deformation seeks to equalize the local and intrinsic variations between the atlas and the subject. Although the variations can be large, the appearance and location of corresponding anatomical structures are still consistent enough for them to be correlated. The intensity difference between spatially corresponding voxels can act as a deforming force, which shifts atlas voxels spatially towards their counterparts in the subject.

An intuitive way to represent the deformation is a 3-D displacement for each voxel in the atlas volume. This would allow each voxel to deform independently, but it will only succeed when the voxels' initial positions are close to their desired positions. It is clear from Figure 6 that the intrinsic variations between individuals make the similarity transformation unable to provide a precise enough initial alignment.

The system employs a more robust representation, a 3-D control grid which is a coarser grid than the voxel grid. The vertices of the control grid are control points that can shift independently in 3-D space. The 3-D displacements of the control points determine the displacements of the voxels they enclose. Figure 7 is a 2-D illustration of the 3-D smooth deformation process. Szeliski used a similar representation in 2-D registration [7] and 3-D registration of surfaces [6]. Since the number of control points is orders of magnitude lower than the number of voxels, there are fewer parameters to estimate, making the deformation process more stable.

Similar to section 2.1, the sum of squared differences (SSD) between corresponding voxels' intensities in the atlas and the subject acts as the error function, and the Levenberg-Marquardt iterative non-linear optimization algorithm is used to determine the smooth deformation parameters. When the change in location of each control point drops below 1 pixel, the smooth deformation is considered to be recovered. To avoid problems with local minima and to improve efficiency, the system uses multi-resolution control grids and multi-resolution image volumes. The control grid

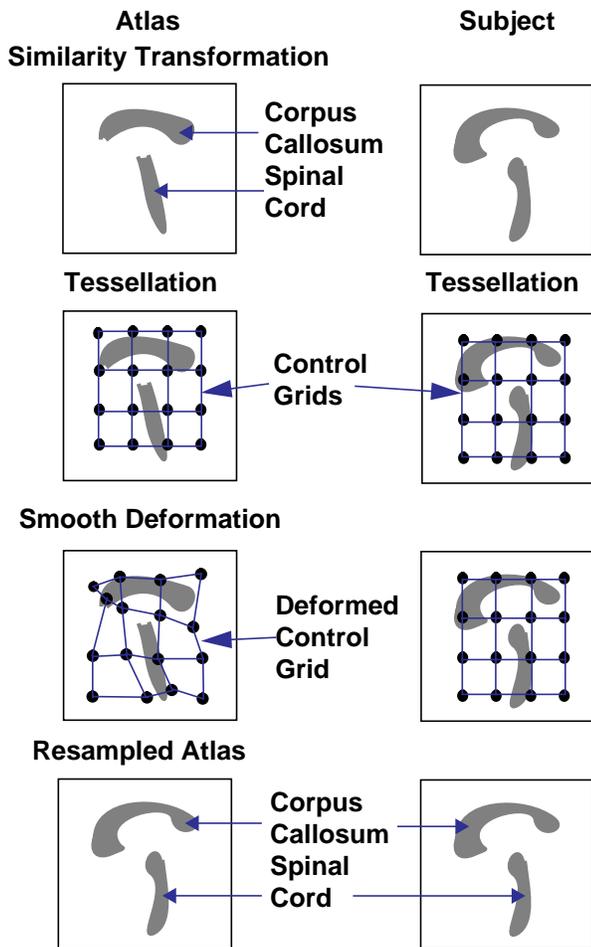


Figure 7: A 2-D illustration of the 3-D smooth deformation. The control points in the atlas shift to match their counterparts in the subject.

has 5 resolution levels, in which the numbers of control points along the x, y, and z directions are $2 \times 2 \times 2$, $3 \times 3 \times 3$, $4 \times 4 \times 4$, $5 \times 5 \times 5$, and $6 \times 6 \times 6$.

Figure 8 shows the result after the smooth deformation is applied to the intermediate result in Figure 6. The atlas is warped in 3-D to match the subject volume. Note that the outlines of the anatomical structures in the deformed atlas roughly align with their counterparts in the subject. The shapes of the lateral ventricles in the atlas appear to be different from those in Figure 6. This is because the atlas has deformed in 3-D, and we are looking at the same cross-section with respect to the subject's data.

2.3. Fine-Tuning Deformation

While the smooth deformation significantly improved the segmentation of individual anatomical structures, there still exists considerable inaccuracy. This is because only the control points can shift independently in 3-D space, and variations smaller than the size of a control grid cell cannot be

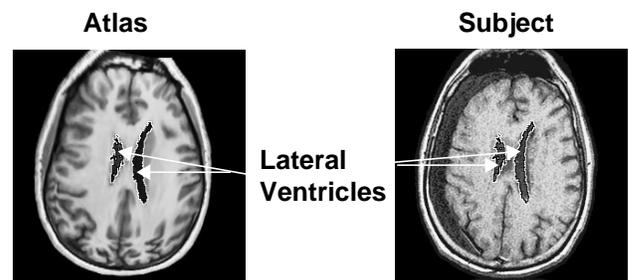


Figure 8: Corresponding slices of the atlas and the subject after the smooth deformation. The atlas is deformed to match the subject. The outlines of the atlas' lateral ventricles approximately align with those in the subject.

accounted for. A fine-tuning deformation is necessary to adjust to the finer details.

To fine-tune the deformation, the intensity difference between spatially corresponding voxels in the atlas and the subject's data again serves as the deforming force. The deformation parameters are the 3-D displacements of each voxel. This is similar to the approach Thirion discussed in [10]. Because the number of deformation parameters is 3 times the number of voxels, the problem is under-constrained. After each iteration, 3-D Gaussian smoothing is applied to the deformation parameters to regularize the problem. This process is similar to an optical flow algorithm. The iteration stops when the root-mean-square (RMS) between the intensities of spatially corresponding voxels decreases by less than 0.5%.

Figure 9 shows the result after applying the fine-tuning deformation to the result in Figure 8. The atlas is further warped in 3-D to better align with the subject. Outlines of anatomical structures in the atlas match well with their counterparts in the subject's data. Note that the outline of the skull is jagged, and the skull is thicker at certain locations. This is because the subject has a pathology that is not present in the atlas. The surrounding structures in the atlas were warped to match with the pathology. This will be discussed more in Section 5.

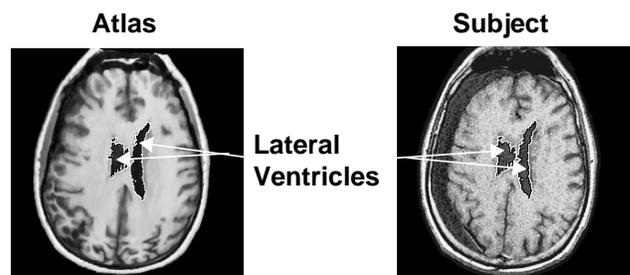


Figure 9: Corresponding slices of the atlas and the subject after the fine-tuning deformation. The outlines of the atlas' lateral ventricles align well with those in the subject.

2.4. Intensity Equalization

The quality of matching is represented by the intensity difference between corresponding voxels in the atlas and the subject's data. However, in general the intensity distribution in the atlas is different from that in the subject. Equalization is necessary to make the intensities comparable.

At each stage in the deformation hierarchy, the system applies a normalization scheme to compensate for the intensity discrepancy:

- Before the similarity transformation, the atlas and the subject volume can be of significantly different orientations and scales. The system normalizes the intensities in the volumes to have the same mean and standard deviation.

- After the similarity transformation, the two volumes are globally aligned, but there exists considerable mismatches between individual structures. The system computes the mean intensity difference between corresponding voxels in the atlas and the subject, and uses the difference to adjust the intensities in the atlas during the smooth deformation process.

- After the smooth deformation, individual structures in the atlas approximately align with those in the subject, especially the ones with distinct shapes and densities, such as the corpus callosum and the skull. The system equalizes the intensities based on this intermediate segmentation. The process involves first computing the intensity histograms of the corpus callosum and the skull in the atlas, and finding the *representative* intensities at the histogram peaks. Similarly, histograms of the regions labelled as corpus callosum and skull in the subject's data are computed. The peak with the highest intensity in the corpus callosum's histogram is considered representative of that structure's intensity, as is the lowest intensity peak in the skull's histogram. The system then linearly adjusts the intensity distribution in the atlas to match these representative intensities. Figure 10 shows examples of the aforementioned histograms.

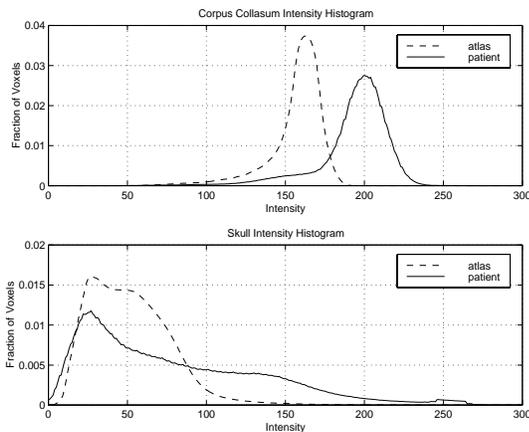


Figure 10: Histograms of the corpus callosum (top) in the atlas (dotted line) and the subject (solid line), as well as those of the skull (bottom).

3. Performance Evaluation

Figure 11 is a close-up of the previous subject's lateral ventricles, overlaid with the segmentation given after each level of the deformation hierarchy. The alignment accuracy increases along the deformation hierarchy.

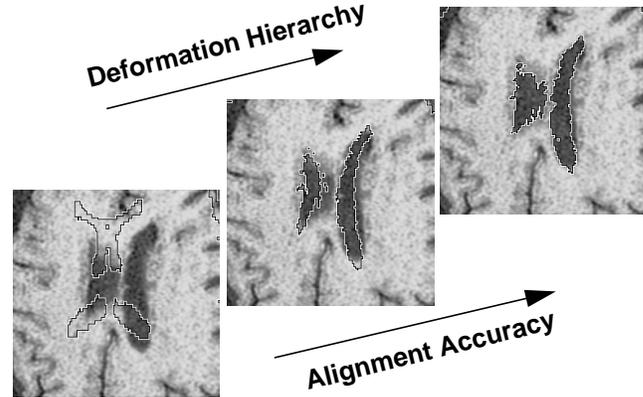


Figure 11: A close-up on the segmentation of the subject's lateral ventricles at each level in the deformation hierarchy.

For quantitative evaluation, we hand segmented the *corpus callosum* in the mid-sagittal plane (a reference plane used by doctors) of 18 subjects, and compared the results with the segmentation given by the customized atlas, as illustrated in Figure 12. We use the percentage of mislabelled voxels relative to the size of the corpus callosum in the hand-segmentation to quantify the segmentation accuracy. Mislabelled voxels include those labelled as corpus callosum in the customized atlas but not in the hand-segmentation and vice versa. For the particular subject shown in Figure 12, 12.3% of the voxels are mislabelled.

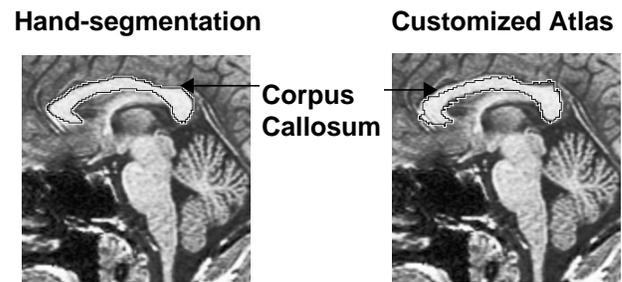


Figure 12: Comparison of segmentation results.

Figure 13 shows distributions of the percentage of mislabelled voxels after each stage of the hierarchical deformation. After the similarity transformation, only 10% of all the cases have less than 20% mislabelled voxels, and 5% of the cases have almost all of the labels wrong. This improved after the smooth deformation, when 61% of the cases have less than 20% mislabelled voxels, the worst 11% of the cases

have 50% mislabelled voxels. After the fine-tuning deformation, 78% of the cases have less than 20% mislabelled voxels, only 11% of the cases have as much as 40% mislabelled voxels. This is consistent with the qualitative observation, that the segmentation accuracy increases after each level of the deformation hierarchy.

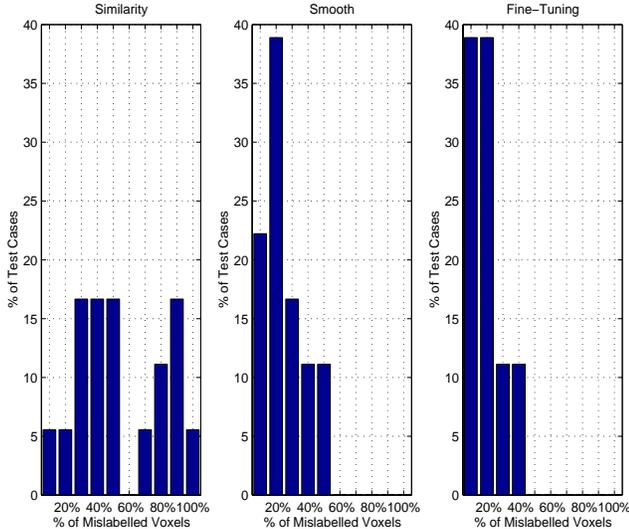


Figure 13: Distributions of the percentage of mislabelled voxels after each stage of the hierarchical deformation.

To show the importance of the intensity equalization procedure, we also computed the percentage of mislabelled voxels of the same 18 subjects, using only the first intensity adjustment (matching the mean and standard deviation of the intensities). As shown in Figure 14, the fraction of cases with less than 20% mislabelled voxels improves from 60% to 78% when the full intensity equalization procedure is used.

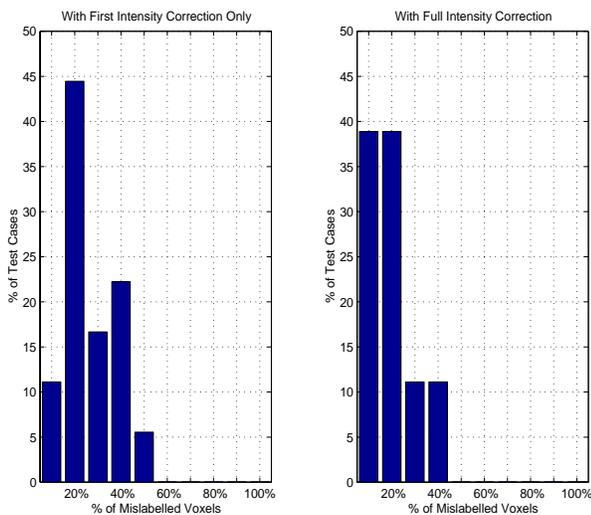


Figure 14: Comparison of the percentage of mislabelled voxels with only initial or the full intensity equalization.

The above evaluations are subject to a single person’s decisions in segmenting the data for comparison. In medical imaging, quantitative evaluation is difficult due to the absence of ground truth. A more objective standard may be developed by fusing multiple experts’ opinions.

4. Statistical Study of Brain Structures

The performance of this system is promising for the quantitative study of the anatomy. This is a fully automatic algorithm. The *similarity transformation* is initialized by trying many random initial transformations and picking the one which gives the minimum SSD [21]. The entire process takes 22 minutes to customize the atlas for a 256x256x124 voxel volume on an SGI computer with three 194 MHz R10000 processors. Parameters such as the number of stochastic samples, the number of levels in the image pyramids, and the number of control points can be tuned to improve efficiency. The convenience and speed makes it possible to study a large number of subjects in great detail. We have used the system to make initial explorations of two areas.

4.1. The Range of Normal Variation

The automated segmentation of anatomical structures enables statistical study of their sizes. Researchers who study abused children examine the sizes of certain brain structures. They spend much time and effort segmenting and assessing the structures for each subject, as mentioned in section 1. Table 1 gives the statistics on these structures’ sizes estimated by our system, based on 105 normal brain images.

Table 1. Statistics on the sizes of brain structures.

Structure Name		mean (mm^3)	standard deviation (mm^3)
3rd ventricle		995	238
4th ventricle		1421	897
amygdala	left	2347	1070
	right	2502	1147
caudate nucleus	left	4885	1120
	right	4620	1059
cerebellum cortex	left	32023	17912
	right	32839	19172
corpus callosum		19758	5354

Table 1. Statistics on the sizes of brain structures.

Structure Name		mean (mm^3)	standard deviation (mm^3)
hippocampus	left	2921	822
	right	2860	753
lateral ventricle	left	6705	1778
	right	6193	1706
midbrain		8211	2220
occipital lobes	left	34632	6693
	right	34507	5659
pons		8900	5740
precentral gyrus	left	13038	2398
	right	17213	3051

This statistical analysis has revealed certain quantitative characteristics of the anatomy. There is considerable variation in the sizes of the same structures across individuals but the variation among normal subjects falls within a certain distribution. The identification and verification of the range of *normal variation* is important in cross-patient analysis and abnormality detection. Experienced doctors acquire a sense of this through practice. Little quantitative study, however, has been done because it requires tremendous effort to manually segment a number of structures for a significant number of subjects. This system brings promise to this area of research.

4.2. Symmetry within the Human Brain

A great number of brain structures have counterparts across the mid-sagittal plane, and they are approximately symmetric (Figure 15, left). Certain pathologies cause *mass effect*, which forces nearby structures to shift from their normal positions and destroys this symmetry (Figure 15, right). Neuroradiologists find this property an important clue in pathology detection. We used this system to compute the absolute difference in volumes of the left and right lateral ventricles for 105 normal brains, and one brain with pathology. In Figure 16, the horizontal axis is the absolute volume difference in mm^3 , the vertical axis is the percentage of normal brains whose lateral ventricle differences fall between certain range. The mean of the absolute volume difference is $794 mm^3$, and the standard deviation is $681 mm^3$. Approximately 72% of the normal brains have lateral ventricles that are symmetric within $1 cm^3$, whereas about 7% of the sub-

jects have lateral ventricles that differ by more than $2 cm^3$. These quantities help reveal the symmetric characteristic of the human brain.

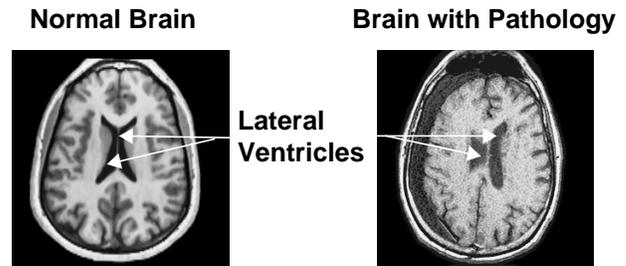


Figure 15: There exists approximate symmetry in the normal brain, but pathologies that cause mass effect destroy this symmetry.

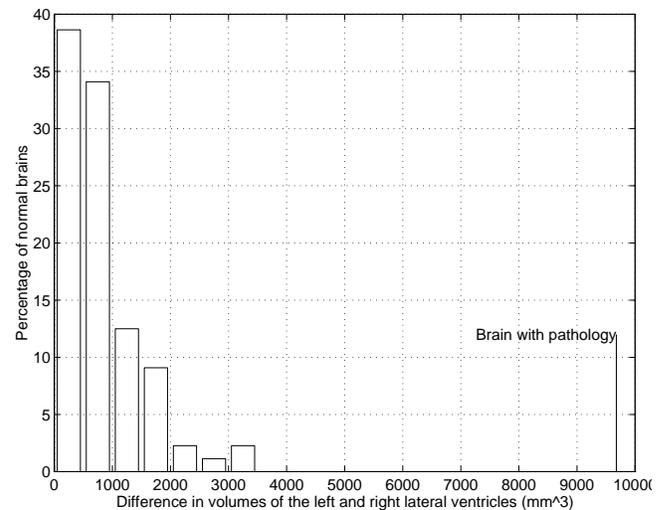


Figure 16: The distribution of the absolute difference in volumes of the left and right lateral ventricles for 105 normal brains.

5. Applications of the Statistics

Quantitative study of anatomical structures brings broad prospects to medical research. It allows measurement of the range of normal variation, and the detection of abnormalities. Statistics on anatomy and pathologies can help express expert knowledge so as to enhance medical education. Moreover, quantitative description of medical image content will facilitate efficient retrieval in the ever increasing medical databases.

For pathology detection, we learned in section 4.2 that the mean of the absolute difference in the volumes of the left and right lateral ventricles in normal brains is $794 mm^3$, and the standard deviation is $681 mm^3$. The pathological case shown in Figure 15, right, has an absolute size difference of $9677 mm^3$, as shown in Figure 16. This is much larger than the normal cases, indicating an abnormal asymmetry, possi-

bly due to mass effect caused by a pathology. The relative size difference between the left and right ventricles indicates the side of the brain containing the pathology.

In section 2.3 the system created the customized atlas for this particular subject. We noted the segmentation of the skull yielded a non-uniform outline. The skull volume estimated by the system is 1035833 mm^3 . Table 1 shows our statistical study of the skull size for 105 normal brains gives a mean of 318079 mm^3 , and a standard deviation of 102558 mm^3 . The size of this subject's labelled skull is beyond the normal range, and hence suggests that an abnormality exists close to the skull.

6. Related Work

The registration of medical images via optimization in transformation space has been an active research area--the comprehensive survey article by van den Elsen et al. [11] lists 161 citations. Two popular schools of registration using image properties are *feature-based* and *voxel-based*. *Feature-based* methods attempt to extract the anatomical structures in different images, and find the correspondences between them. They have the characteristic of being efficient in representation and independent of imaging modality. However, feature-based registration is critically dependent on the quality of the feature extraction, which is not trivial since anatomical structures tend to have complex shapes and ill-defined boundaries. Human interaction is generally necessary to help select and extract features or to guide the matching procedure. Consequently, it is subject to user subjectivity, time-consuming, and inconvenient. As an alternative, *voxel-based* algorithms obviate the need for an explicit segmentation. The most intuitive voxel-based approach uses *voxel intensities*.

Bajcsy et al. developed a system that elastically deforms a 3D atlas to match anatomical brain images [8], [9]. The atlas is modelled as a physical object and is given elastic properties. Although their approach is similar to our current one, they assume the intrinsic variations between people can be modelled by an elastic deformation, whereas we only enforce smooth deformation in an intermediate stage. Without user interaction, their atlas can have difficulty matching complicated object boundaries. Their method is computationally expensive and requires interactive and time-consuming preprocessing.

Christensen et al. presented a method very close to ours, except that they used a fluid dynamic model for the deformation [1], [2]. It takes 1.8 hours to match two $128 \times 128 \times 100$ volumes on a 16384-processor MasPar computer.

In [10], Thirion takes a similar approach to ours, except that he assumes the volumes are already globally aligned, and applies optical flow from the beginning. To reduce computation time, he used the gradient of the subject volume instead of the deformed atlas, because the computation of the

latter is more expensive, requiring tri-linear interpolation of each voxel's gradient. However, this quicker method may cause errors when the deformed atlas does not resemble the subject closely. Because optical flow relies heavily on the constant brightness assumption, it is prone to failure when there are large intensity variations between different images.

Although voxel intensity-based approaches have shown encouraging results, they are problematic when there are intensity inhomogeneities. Moreover, they only work for multi-modal data if there exists a linear mapping between intensity values. Viola and other researchers have investigated registration based on mutual information (MI), [17], [18], [19]. MI is a basic concept from information theory, which measures the statistical dependence between two random variables, or the amount of information that one variable contains about the other. The MI registration criterion assumes that the statistical dependence between corresponding voxel intensities is maximal if the images are geometrically aligned. Because no assumptions are made regarding the nature of this dependence, the MI criterion is highly data independent. Current applications of MI to registration only perform affine transformations to register image data of the same person from different modalities. The possibility of applying the MI criterion in deformable registration remains to be studied.

7. Discussion and Future Directions

We introduced a system that automatically segments and classifies features in 3-D images. The system accomplishes this by starting with an *atlas*, an MRI of a normal brain hand segmented and classified by an expert. Given a subject's data, the atlas is warped using a 3-D *hierarchical deformable matching algorithm* until it closely matches the subject, i.e. the atlas is *customized* for the subject. Then the *customized atlas* can be used directly to segment and classify the subject's anatomical structures.

It takes 22 minutes to segment and classify 144 brain structures in a $256 \times 256 \times 124$ voxel image, while similar data took an expert 8 months to segment and label. We performed quantitative evaluation of the segmentation for one anatomical structure. Of 18 subjects for which classification correctness was examined voxel by voxel, a mislabelling rate of less than 20% was achieved for 14 subjects, and less than 10% for 7 subjects.

The efficiency and accuracy of the system's performance show promise for detailed quantitative studies of brain structures. The system's potential has been demonstrated for finding the range of normal variation of anatomical structures, studying brain symmetry, and detecting abnormalities.

Several issues remain to be addressed. Currently the performance suffers when there are significant differences between the intensity distributions of the atlas and the sub-

ject. A more effective intensity normalization scheme or a criterion not based on intensity similarity may improve the algorithm. More objective and extensive validation of the segmentation accuracy is important to justify the quantitative studies based on the system's results. To study the inherent properties of anatomical structures, statistics of properties other than volume will need to be examined, and larger databases should be collected. Close collaboration with medical professionals will be essential in each of these areas.

Acknowledgments

The authors are thankful to the Brigham and Women's Hospital of the Harvard Medical School for the atlas data, and to Kate Fissell in the Carnegie Mellon Psychology Department for MRI data of 88 normal brains. We are also grateful to Peter Rander and Justin Boyan for their insightful comments.

References

- [1] Christensen et al., "Individualizing Neuroanatomical Atlases Using A Massively Parallel Computer", IEEE Computer, pp. 32-38, January 1996.
- [2] Christensen et al., "Deformable Templates Using Large Deformation Kinematics", IEEE Transactions on Image Processing, September 1996.
- [3] Horn, "Closed-form solution of absolute orientation using unit quaternions", Journal of the Optical Society of America, Vol. 4, No. 4, pp 629-642, April, 1987.
- [4] Lavallee et al., "Image-Guided Operating Robot: A Clinical Application in Stereotactic Neurosurgery", Computer-Integrated Surgery--Technology and Clinical Applications, Taylor et al., pp. 346-351.
- [5] McInerney and Terzopoulos, "Deformable models in medical image analysis: a survey", Medical Image Analysis, Vol. 1, No. 2, pp 91-108, 1996.
- [6] Szeliski, "Matching 3-D Anatomical Surfaces with Non-Rigid Deformations using Octree-Splines", International Journal of Computer Vision, Vol. 18, No. 2, pp 171-186, 1996.
- [7] Szeliski and Coughlan, "Spline-Based Image Registration", Digital Equipment Corporation, Cambridge Research Lab, Technical Report 94/1.
- [8] Bajcsy and Kovacic, "Multiresolution Elastic Matching", Computer Vision, Graphics, and Image Processing, Vol. 46, pp 1-21, 1989.
- [9] Gee, Reivich and Bajcsy, "Elastically Deforming 3D Atlas to Match Anatomical Brain Images", Journal of Computer Assisted Tomography, Vol. 17, No. 2, pp 225-236, 1993.
- [10] Jean-Philippe Thirion, "Fast Non-Rigid Matching of 3D Medical Images", INRIA, Technical Report No. 2547, May, 1995.
- [11] van den Elsen, Pol, and Viergever, "Medical Image Matching--A Review with Classification", IEEE Engineering in Medicine and Biology, Vol. 12, No. 1, pp. 26-39, 1993.
- [12] Antoine Maintz, van den Elsen, and Viergever, "Evaluation of Ridge Seeking Operators for Multimodality Medical Image Matching", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 18, No. 4, April, 1996.
- [13] Szekeley et al., "Segmentation of 2-D and 3-D Objects from MRI Volume Data Using Constrained Elastic Deformations of Flexible Fourier Contour and Surface Models", Medical Image Analysis, Vol. 1, No. 1, pp. 19-34, 1996.
- [14] Sclaroff and Pentland, "On Modal Modelling for Medical Images: Underconstrained Shape Description and Data Compression", M.I.T. Media Laboratory Perceptual Computing Section Technical Report No. 275.
- [15] Sclaroff and Pentland, "Modal Matching for Correspondence and Recognition", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 17, No. 6, June, 1995.
- [16] Moghaddam Nastar and Pentland, "A Bayesian Similarity Measure for Direct Image Matching", M.I.T Media Laboratory Perceptual Computing Section Technical Report No. 393.
- [17] Wells III et al., "Multi-Modal Volume Registration by Maximization of Mutual Information", Medical Image Analysis, Vol. 1, No. 1, pp 35-51, 1996.
- [18] Viola and Wells III, "Alignment by Maximization of Mutual Information", Proceedings of the fifth International Conference on Computer Vision, pp 16-23, 1995.
- [19] Maes et al., "Multimodality Image Registration by Maximization of Mutual Information", IEEE Transactions on Medical Imaging, Vol. 16, No. 2, April, 1997.
- [20] Press et al., "Numerical Recipes in C", Cambridge University Press, 1992.
- [21] Pokrandt, "Fast Non-Supervised Matching: A Probabilistic Approach", Karlsruhe internal technical report, 1996.
- [22] Mei Chen, Takeo Kanade, Henry A. Rowley, Dean Pomerleau, "Anomaly Detection through Registration", Image Registration Workshop, NASA Goddard Flight Center, November, 1997.