Medical Image - Atlas Registration Using Deformable Models for Anomaly Detection

Mei Chen **Takeo Kanade** **Dean Pomerleau**

{meichen, tk, pomerlea}@cs.cmu.edu
School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract[1]

We introduce a system that automatically segments and classifies structures in brain MRI volumes. It segments 144 structures of a 256x256x124 voxel image in 18 minutes on an SGI computer with four 194 MHz R10K processors. The algorithm uses an atlas, a hand segmented and classified MRI of a normal brain, which is warped in 3-D using a hierarchical deformable registration algorithm until it closely matches the subject. This customized atlas contains the segmentation and classification of the subject's anatomical structures. The system has processed 198 MRIs of normal brains, and 3 MRIs and 1 CT of brains with pathologies. Quantitative evaluations yield high segmentation accuracy. Combined with domain knowledge, the registration algorithm is able of detecting asymmetries and abnormal variations in the subject's data that indicate the existence and location of pathologies.

1. Anomaly Detection

Medical experts detect anomalies by comparing a particular subject to normal cases. In Figure 1, the normal brain is approximately symmetric across the center line, e.g. the pair of lateral ventricles have symmetric shapes and sizes. However, the pathological brain lost this symmetry, which is an indication of the existence and location of an anomaly, i.e. a bleeding.

1. This research has been sponsored in full or in part by Office of Naval Research (ONR) under Contract N00014-95-1-0591. 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 ONR or the U.S. Government.

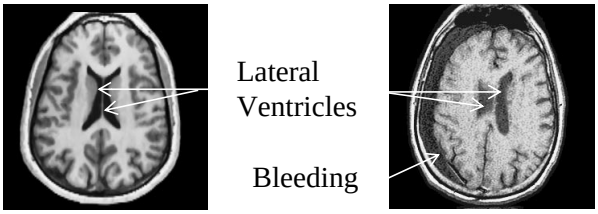

Figure 1. Compare a normal (left) and an abnormal brain.

Our research assists the diagnosis by automating the comparison process. Our reference is an *atlas*, a magnetic resonance imaging (MRI) volume of a normal brain. It contains 123 slices, and each slice is a 256x256 pixel matrix. An expert spent 8 months to manually segment 144 anatomical structures, and gave each structure a unique label. The labels of the structures were color-coded to illustrate the segmentation, Figure 2. We developed a 3-D *hierarchical deformable registration algorithm*, which constructs a mapping from the atlas to a subject. By applying the mapping to the color-coded segmentation of the atlas, we acquire a *customized* segmentation of the subject's anatomical structures, Figure 2. We combine the registration result and the segmentation information with domain knowledge to aid anomaly detection.

2. Problem definition

The data we work with is T1 weighted MRI. An MRI volume is a series of parallel cross-sections along one of three principal axes, see Figure 4.

The atlas may differ from a subject in two ways. Genetic and environmental factors cause variations in the shape, size, density, and location of anatomical structures, i.e. variations that are local and intrinsic. The lack of standards in the data acquisi-

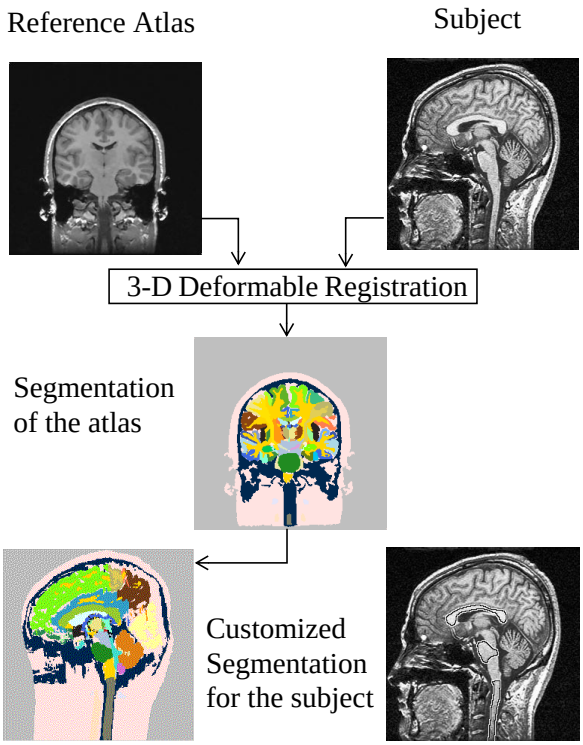

Figure 2. Illustration of the comparison process.

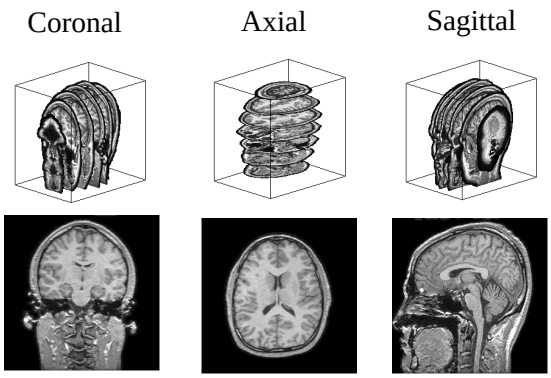

Figure 3. MRI scans along the principal axes.

tion process causes variations in the orientation, scale, resolution, and intensity consistency, i.e. variations that are global and extrinsic. The registration process compensating both kinds of variations to match the atlas to the subject.

3. Hierarchical Deformable Registration

Because variations can be geometrical or intensity-wise, intrinsic or extrinsic, we adopt a hierarchical approach. It involves two interleaving hierarchies, one is a hierarchy of intensity equalization schemes, one is a hierarchy of geometrical deformable models, see Figure 4. First, we grossly equalize the intensities of the atlas and the subject volumes, and transform the atlas to globally align with the subject (section 3.1). Based on this initial alignment, we apply a more localized intensity normalization, and employ a smooth deformation to roughly match the anatomical structures of the atlas to those in the subject (section 3.2). This correspondence then allows a more informed intensity transformation between the two volumes, and a fine-tuning deformation fits the atlas more precisely to the subject. (section 3.3). This automatic algorithm starts by trying random similarity transformations and picking the optimum one[8]. Iterative optimization algorithms are used to determine the deformation parameters at all levels.

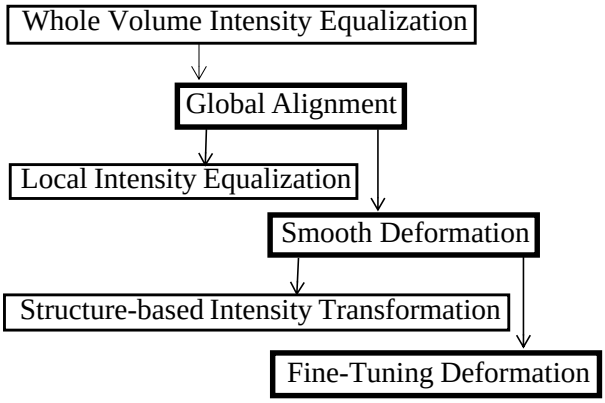

Figure 4. Hierarchical Deformable Registration.

3.1. Global Alignment

The first level deformable model adjusts the extrinsic variations between the atlas and the subject by rotating, translating, and uniformly scaling the atlas to match the subject. The quality of a match is measured by the sum of squared differences (SSD) between the intensities of corresponding voxels in the two volumes. Separate imaging processes may result in different intensity distributions in the atlas and the subject volumes. The first level intensity equalization levels the mean and standard deviation in the two volumes to roughly correct this discrepancy. The Levenberg-Marquardt non-linear optimization algorithm is used to iteratively adjust the transformation to reduce the SSD. We use multi-resolution processing and stochastic sampling for efficiency and to help prevent the optimization from being trapped in local minima.

Figure 5 shows example cross-sections of the atlas and a subject, before and after the global alignment. The atlas is rotated, scaled, and translated to grossly match the subject. The label of an anatomical structure, the corpus callosum, in the transformed atlas is directly applied to the subject data. It does not align well with its counterpart in the subject.

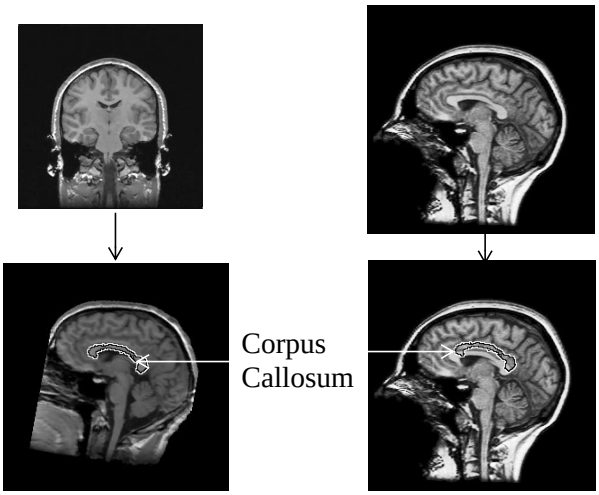

Figure 5. The atlas (left) and the subject, before and after the global alignment.

3.2. Smooth Deformation

The local and intrinsic variations between the atlas and the subject can be large, but humans can still correlate corresponding structures based on similarities in shapes, intensities, and locations. Therefore, the intensity difference between corresponding voxels in the atlas and the subject can act as a deforming force, which shifts atlas voxels spatially towards their counterparts in the subject volume. To make the intensities comparable, the second level intensity equalization evens the mean and standard deviation of the overlapping portion of the atlas and subject volumes.

Intuitively we may allow each atlas voxel to deform freely. However, to avoid local minima, the voxel's initial position must be close to its desired position. It is clear from Figure 5 that the intrinsic variations between individuals make the global alignment unable to provide a precise enough initialization. A more robust representation is a 3-D control grid that is coarser than the voxel grid. The vertices of the control grid are control points that shift independently in 3-D. Their 3-D displacements determine the displacements of the voxels they enclose. Figure 6 is a 2-D illustration of the second level deformable model in 3-D [3], [4]. We use control grids in multi-resolution to improve efficiency and avoid local minima. Since the number of control points is orders of magnitude fewer than the number of voxels, there are fewer parameters to estimate, making the deformation process more stable.

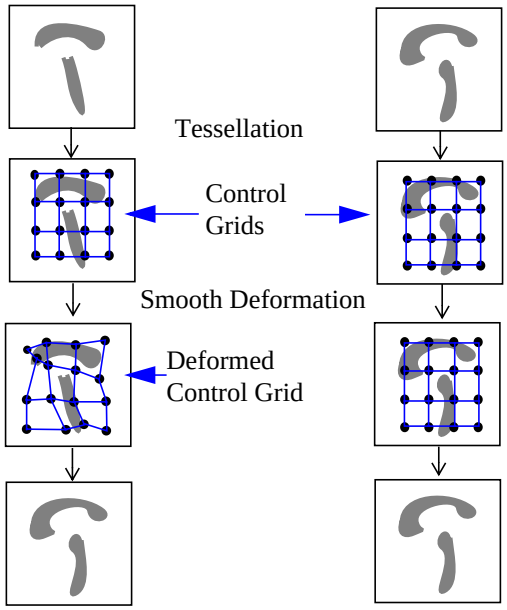

Figure 6. A 2-D illustration of the 3-D smooth deformation. The control points in the atlas (left) shift to match their counterparts in the subject.

Figure 7 shows the effect of applying the smooth deformation to the intermediate result in Figure 5. The atlas is warped in 3-D to match the subject. The outline of the corpus callosum in the deformed atlas roughly align with its counterpart in the subject.

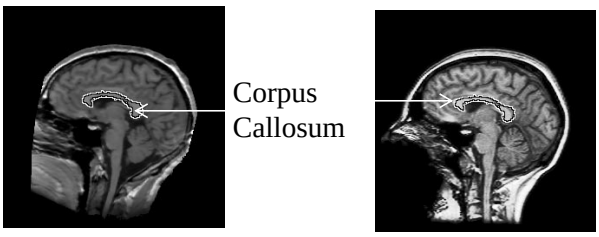

Figure 7. The atlas (left) and the subject, after the smooth deformation.

3.3. Fine-tuning deformation

Smooth deformation only allows the control points to shift freely in 3-D, variations smaller than the size of a control grid cell cannot be accounted for. The last level deformable model fine-tunes the registration by permitting each atlas voxel to shift independently in 3-D [6]. The problem is underconstrained because the number of deformation parameters are the voxels' 3-D displacements, which is 3 times the number of voxels. 3-D Gaussian smoothing is applied to regularize the deformation. The intensity difference between spatially corresponding voxels again serves as the deforming force. Since the smooth deformation roughly registers individual structures in the atlas with those in the subject, the distributions of a high intensity (corpus callosum) and a low intensity (skull) structure can jointly determine a linear transformation that further equalizes the two volumes' intensity distributions.

Figure 8 shows the result of applying the fine-tuning deformation to the result in Figure 7. The outline of the corpus callosum in the deformed atlas match well with its counterpart in the subject.

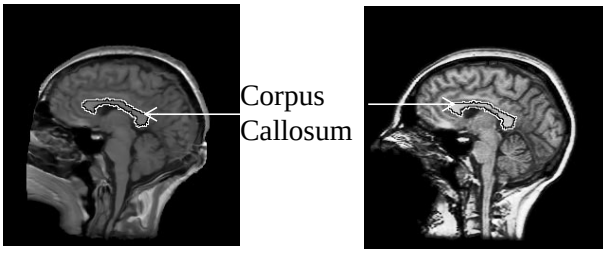

Figure 8. The atlas (left) and the subject, after the fine-tuning deformation.

Figure 9 is another example of the hierarchical deformable registration. Outlines of the lateral ventricles in the deformed atlas are applied to the subject volume to show the improvement of segmentation accuracy.

4. Performance Evaluation

. The entire algorithm takes 18 minutes to customize the atlas segmentation for a 256x256x124 voxel image on an SGI computer with four 194 MHz R10K processors. Parameters can be tuned to further improve the efficiency.

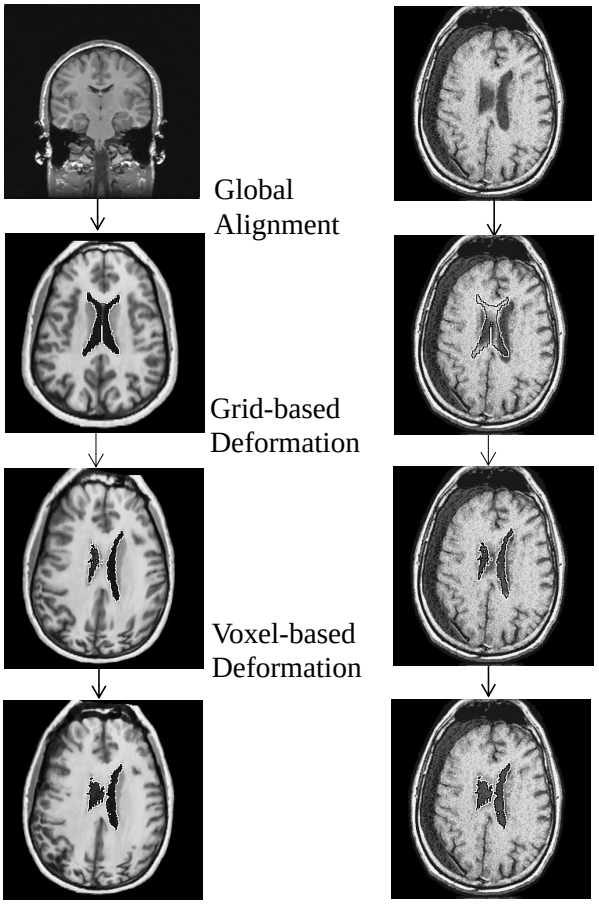

Figure 9. Progressive results of the hierarchical registration between the atlas (left) and a subject.

We processed 198 MRIs of normal brains, and 3 MRIs and 1 CT of brains with pathologies. The number of voxels ranges from 0.8 to 16 million. To evaluate the accuracy, we hand segmented the corpus callosum in the mid-sagittal plane (a reference plane used by doctors) of 42 subjects, and compared them with the segmentations given by the algorithm. We use the number 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 by the algorithm but not in the hand-segmentation and vice versa. Figure 10 shows distributions of the percentage of mislabelled voxels after each level of the hierarchical deformation. After the global alignment only 14% of the cases have less than 20% mislabelled voxels. The smooth deformation improved this ratio to 43%, while 14% of the cases have less than 10% mislabelled voxels. The final fine-tuning brought the ratio of mislabelled voxels to below 20% for all

cases, and to below 5% for 90% of the cases. This is consistent with the qualitative observation in Figure 9.

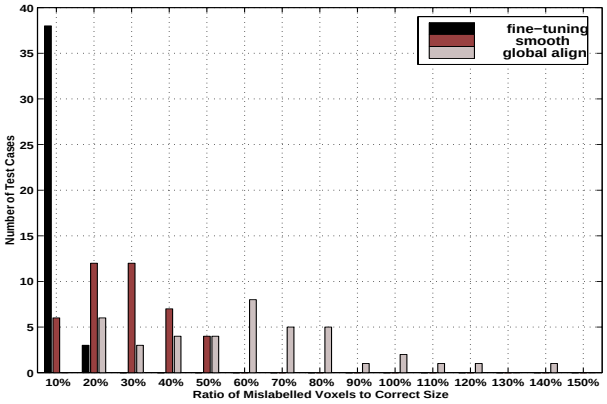

Figure 10. Distribution of the ratio of mislabelled voxels after each stage of the deformation.

5. Anomaly detection through registration

Our hierarchical deformable registration algorithm has proven effective at matching and segmenting anatomical structures. It also has potential in detecting and locating pathologies.

5.1. Brain Symmetry

As mentioned in section 1, human brain has approximate symmetry across the mid-sagittal plane. Certain pathologies cause *mass effect*, which forces nearby structures to shift from their normal positions and destroys this symmetry. Doctors find the lack of symmetry an indication of a pathological condition. We detect asymmetry by matching a subject's mirror volume (with left and right flipped) to the original volume. For a normal brain, little deformation is needed to align its mirror volume to itself; for a brain with pathologies that cause mass effect, the deformation will be significant. The direction of the deformation suggests the location of the anomaly. Figure 11 shows results on three subjects of different conditions.

The first row shows an axial cross-section of each subject's volume. The second row is the cross-section overlaid with the deformation flows for matching the subject's mirror volume to the original volume. For the normal brain, the magnitude of the deformation vectors is negligible because it is

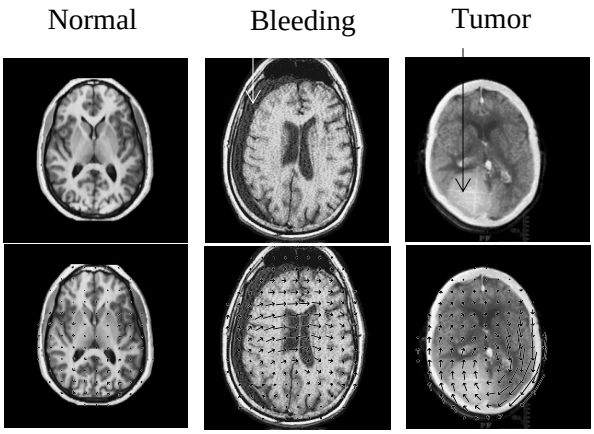

Figure 11. Deformation flows of matching the mirror volume to the original volume.

approximately symmetric. For the brain with a bleeding, the deformation flows have large magnitudes, and exhibit a uniform direction. This direction suggests the source of the mass effect. For the brain with a tumor, the magnitudes of the deformation flows are large, and the directions has a swirl pattern. This is because the tumor is so close to the central line that its "counterpart" in the mirror volume partly matched with itself. This method is effective in detecting anomalies that affect brain symmetry.

5.2. Detect Abnormal Variations

Despite the large range of intrinsic variations between individuals, there exists a distinction between the normal range of natural shape variability and pathology-afflicted changes. Figure 12 compares the skull label for a normal brain and a brain with pathology. The skull of the normal brain has uniform thickness, with variations within a small range, whereas the thickness of the skull in the pathological brain varies significantly. Since this is beyond the normal variation of skull thickness, it may be an indication of an anomaly.

6. Related Work

Van Den Elsen et al. [7] did a comprehensive survey of medical image registration via optimization in transformation space. Bajcsy et al. [5] elastically deform a 3D atlas to match brain images. The method requires user interaction and is computationally expensive. Christensen et al. [1] used a fluid dynamic deformable model. It takes 1.8 hours

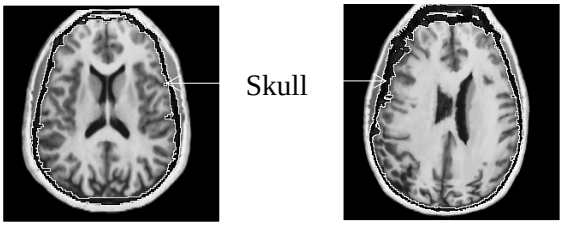

Figure 12. Comparing the *skull* label between a normal brain and a brain with pathology.

to match a 1.6 million voxel image on a 16384-processor MasPar. Thirion [6] assumes the volumes are already globally aligned, and uses optical flow for deformation. Since optical flow relies heavily on the constant brightness assumption, it is sensitive if the intensity variations or deformations are large. Viola et al. [8], Maes et al.[9] investigated registration based on mutual information. Currently it can only determine affine transformations between images of the same person. Cross-subject registration is being studied.

7. Conclusion and Discussion

We have introduced a hierarchical deformable registration algorithm that takes minutes to automatically segment and classify features in brain MRI volumes. By employing the symmetry characteristic of normal brains, it can detect and locate anomalies that affect symmetry; by applying knowledge of normal variations in anatomy, it may be able to find pathologies that incur shape, size, or density changes. This algorithm's speed allows us to examine a significant number of subjects, so as to characterize normal variations in the shape, size, and location of anatomical structures. This characterization can then be used to improve the segmentation accuracy, by guiding the deformation process and limiting the space of deformations that must be searched; on the other hand, it can also aid anomaly detection, by detecting the abnormal deformations required to match an anomalous structure to the normal atlas.

Acknowledgments

The authors are thankful to the Brigham and Women's Hospital of the Harvard Medical School for the brain atlas. We are grateful to Carnegie Mellon Psychology Department, Western Psychiatric Institute and Clinic, and National Institute of Health for the MRI data. We are also beholden to our colleagues for their comments.

References

[1] Christensen et al., "Individualizing Neuroanatomical Atlases Using A Massively Parallel Computer", IEEE Computer, pp. 32-38, January 1996.

[2] McInerney and Terzopoulos, "Deformable models in medical image analysis: a survey", Medical Image Analysis, Vol. 1, No. 2, pp 91-108, 1996.

[3] 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.

[4] Szeliski and Coughlan, "Spline-Based Image Registration", Digital Equipment Corporation, Cambridge Research Lab, Technical Report 94/1.

[5] 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.

[6] Jean-Philippe Thirion, "Fast Non-Rigid Matching of 3D Medical Images", INRIA, Technical Report No. 2547, May, 1995.

[7] 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.

[8] Viola and Wells III, "Alignment by Maximization of Mutual Information", Proceedings of the fifth International Conference on Computer Vision, pp 16-23, 1995.

[9] Maes et al., "Multimodality Image Registration by Maximization of Mutual Information", IEEE Transactions on Medical Imaging, Vol. 16, No. 2, April, 1997.

[10] Pokrandt, "Fast Non-Supervised Matching: A Probabilistic Approach", Karlsruhe internal technical report, 1996.

[11] Mei Chen et al., "Anomaly Detection through Registration", Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp 304-310, June, 1998.