

# 3-D Deformable Registration of Medical Images Using a Statistical Atlas

Mei Chen<sup>1</sup>, Takeo Kanade<sup>1</sup>, Dean Pomerleau<sup>1,2</sup>, Jeff Schneider<sup>1</sup>

<sup>1</sup>Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213

{meichen, tk, pomerlea, schneide}@cs.cmu.edu

[www.cs.cmu.edu/~meichen/registration.html](http://www.cs.cmu.edu/~meichen/registration.html)

<sup>2</sup>AssistWare Technology, 123000 Perry Highway, Wexford, PA 15090

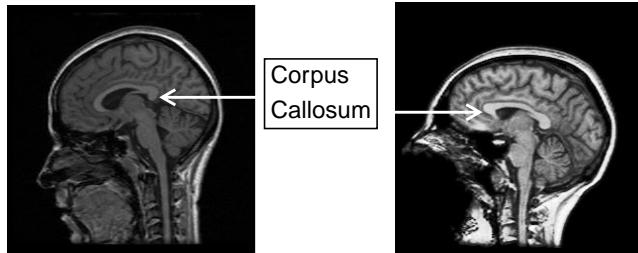
## Abstract

Registration between 3-D images of human anatomies enables cross-subject diagnosis. However, innate differences in the appearance and location of anatomical structures between individuals make accurate registration difficult. We characterize such anatomical variations to achieve accurate registration.

We represent anatomical variations in the form of statistical models, and embed these statistics into a 3-D digital brain atlas which we use as a reference. These models are built by registering a training set of brain MRI volumes with the atlas. This associates each voxel in the atlas with multi-dimensional distributions of variations in intensity and geometry of the training set. We evaluate statistical properties of these distributions to build a statistical atlas. When we register the statistical atlas with a particular subject, the embedded statistics function as prior knowledge to guide the deformation process. This allows the deformation to tolerate variations between individuals while retaining discrimination between different structures. This method gives an overall voxel mis-classification rate of 2.9% on 40 test cases; this is a 34% error reduction over the performance of our previous algorithm without using anatomical knowledge. Besides achieving accurate registration, statistical models of anatomical variations also enable quantitative study of anatomical differences between populations.

## 1. Motivation

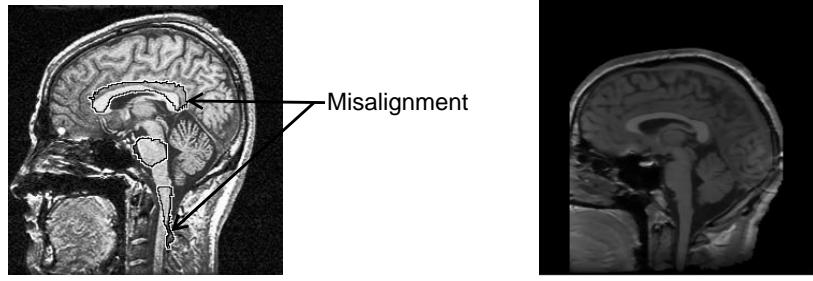
Registration between volumetric images of human bodies enables cross-subject diagnosis and post-treatment analysis. However, due to genetic and life-style factors, there are innate variations among individuals in the appearance and location of anatomical structures. Figure 1 displays cross-sections of T1-weighted magnetic resonance imaging (MRI) volumes of two non-pathological brains. The example structure, corpus callosum, has different intensity, shape, size, and location in these two brains. For registration algorithms that use only intensity or shape templates to achieve correspondence, results are typically poor due to these inherent variations.



**Figure 1.** Innate variations between individuals.

Currently there exist many intensity correspondence based registration algorithms [1], [3], [5], [13]. Figure 2 shows a registration result using method [5]. The right image volume is deformed to register with the left image volume, and outlines of its anatomical structures are overlaid on the left image to illustrate the alignment. Note that there is significant misalignment between the deformed structures and the real structures. This is because the shape and density of anatomical structures in the two volumes are considerably different, and a method using intensity correspondence cannot address the difference.

Knowledge of anatomical variations provides information that can guide the registration process and improve accuracy. Characterization of such variations also facilitates quantitative study of anatomical differences between populations, as well as anomaly detection. We capture and model non-pathological anatomical differences between individuals, and use this knowledge to achieve accurate registration.



**Figure 2.** The right image volume is deformed to register with the left one. Outlines of several deformed anatomical structures are overlaid on the real structures

## 2. Problem Definition

We collected 105 T1-weighted MRI volumes of non-pathological brains, and use them as the training set for knowledge extraction. Examples from the training set are shown in the top row of Figure 3. Apart from the intrinsic differences between different people's brain structures, there are also extrinsic differences in image orientation, scale, resolution, and intensity.

In order to capture the anatomical variations in the training set, we compare each MRI volume to a common reference. Our reference is a 3-D digital atlas, which is a T1-weighted MRI of a non-pathological brain, accompanied by expert classification of its anatomical structures. Note that, this atlas is an example of a normal brain, not an average brain of a population. The method for comparison is an automatic 3-D deformable registration algorithm that was previously developed<sup>[13]</sup>. This method first eliminates the extrinsic variations between image volumes with preprocessing, then extracts the intrinsic anatomical variations by finding the deformable mapping between each image volume and the atlas.

The intrinsic variations are abstracted into a computational model. During registration, this model functions as prior knowledge to guide the registration process to tolerate anatomical variations, and to achieve higher accuracy.

## 3. Capturing Anatomical Variations

Different image acquisition processes result in variations in the 3-D orientation, position, resolution and intensity of image volumes in the training set. Differences in head size also add variation in the scales of the image volumes. These variations are extrinsic to the anatomical variabilities, and thus need to be removed before the intrinsic variations can be extracted.

### 3.1. Eliminating Extrinsic Variations

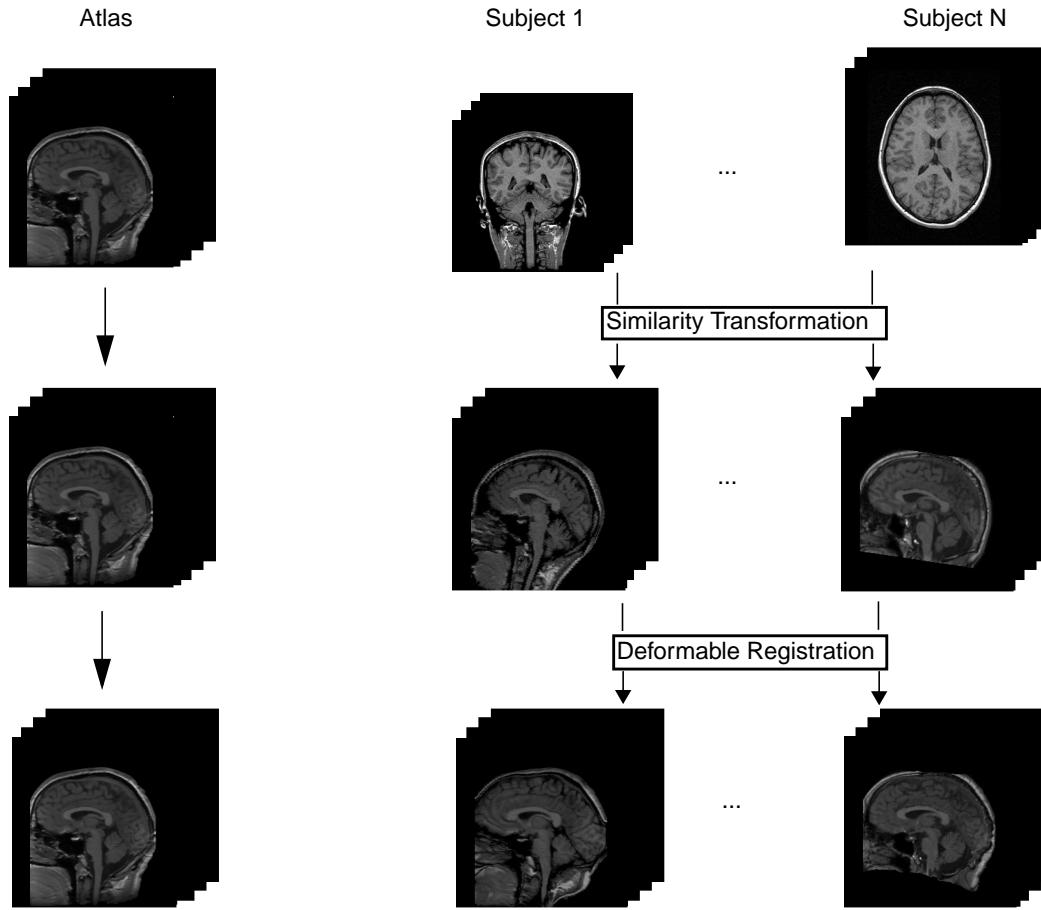
The registration algorithm employed for atlas-training set comparison<sup>[13]</sup> consists of a hierarchy of deformable models, of which the first level is a similarity transformation, which addresses the extrinsic geometric variations between different subject volumes via 3-D rotation, scaling, and translation. As a result, each subject volume in the training set has roughly the same orientation, size, and location as that of the atlas. The transformed subject volume is resampled to match the resolution of the atlas. A multi-level intensity equalization scheme is interwoven into the deformation hierarchy to adjust the differences in intensity distributions. The middle row in Figure 3 shows the result of having removed the extrinsic differences between the atlas and the training set.

### 3.2. Extracting Intrinsic Anatomical Variations

After the removal of extrinsic variations, intrinsic variations are apparent as the misalignment between anatomical structures in the subject volumes and the atlas. The employed registration algorithm captures this information by aligning the corresponding structures through 3-D deformation, and recording the 3-D displacement, as shown in the last row of Figure 3. Therefore, after aligning each subject's anatomical structures with those in the atlas, each atlas voxel is associated with two distributions: one is an intensity distribution of corresponding voxels in the subject volumes; the other is a geometric distribution of the 3-D displacement between the atlas voxel and the corresponding voxels in the subject volumes. The former contains density variations of anatomical structures over a population (density is reflected in image intensity), while the latter embodies geometrical variations of these structures, such as shape, size, and location.

## 4. Modeling Anatomical Variations

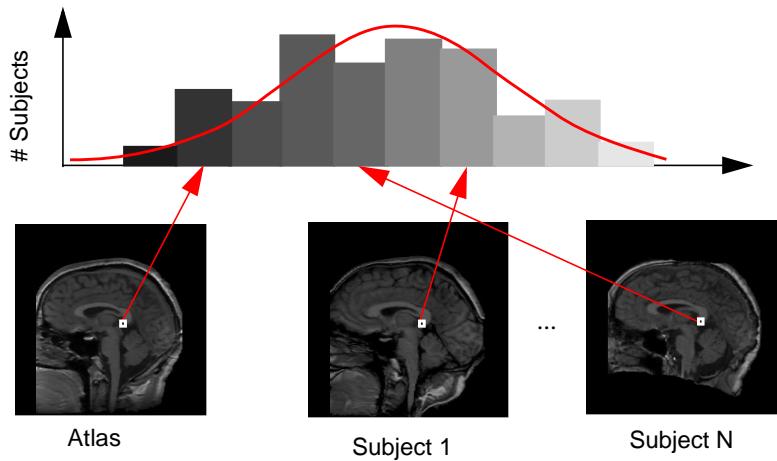
The purpose of capturing anatomical variations is to achieve accurate registration. We characterize these variations in a statistical manner, so as to employ them as prior knowledge in statistical models.



**Figure 3.** The process of removing extrinsic variations and extracting intrinsic variations.

#### 4.1. Modeling Density Variations

Once the training set is deformed to register with the atlas, each atlas voxel corresponds with its counterpart in each of the subject volumes. The histogram of their intensities captures tissue density variations in a population (Figure 4).



**Figure 4.** Model density variations at each atlas voxel as a 1-D distribution.

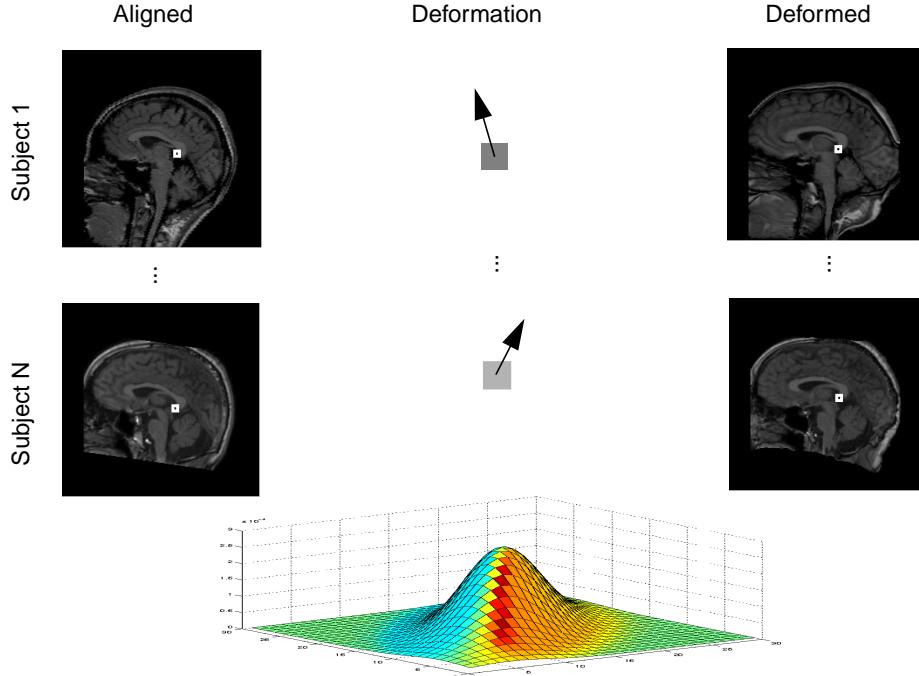
The intensity histogram at each atlas voxel is modeled as a 1-D Gaussian distribution,  $P\langle dI|D \rangle$ :

$$P\langle dI|D \rangle = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(dI - \mu)^2}{2\sigma^2}} \quad (1)$$

where  $dI = I_s - I_a$ , while  $I_s$  and  $I_a$  are corresponding voxel intensities in the subject volume and the atlas.  $D$  is the 3-D deformation between them.  $\mu$  is the mean intensity difference between the training set and the atlas at this voxel;  $\sigma^2$  is the variance of the intensity difference distribution.  $dI$  has been adjusted for intensity variations caused by image acquisition processes.

#### 4.2. Modeling Geometric Variations

After the training set is deformed to register with the atlas, the 3-D displacements between each atlas voxel and its counterparts in the training set embody the geometric variations between individuals. The distribution of the variations can be captured in a 3-D histogram. Figure 5 shows a 2-D illustration.



**Figure 5.** Model geometric variations at each atlas voxel as a 3-D distribution.

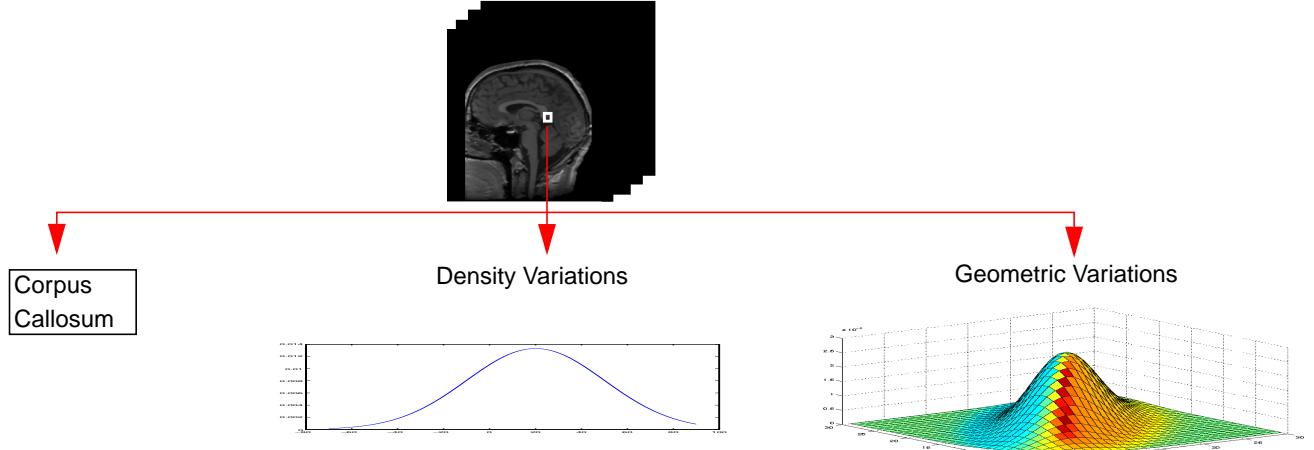
The 3-D histogram of displacements at each atlas voxel is modeled as a 3-D Gaussian distribution,  $P(D)$ :

$$P(D) = \frac{1}{\sqrt{(2\pi)^3 |\Phi|}} e^{-\frac{(\bar{\Delta}\vec{\vartheta} - \vec{\omega})^T \Phi^{-1} (\bar{\Delta}\vec{\vartheta} - \vec{\omega})}{2}} \quad (2)$$

here  $\bar{\Delta}\vec{\vartheta}$  is the 3-D displacement between the atlas voxel and its counterparts in the training set,  $\vec{\omega}$  is the mean 3-D displacement at this atlas voxel, and  $\Phi$  is the 3x3 covariance matrix of the distribution.  $\bar{\Delta}\vec{\vartheta}$  has been adjusted for extrinsic variations.

#### 4.3. A Statistical Atlas

The original atlas was one particular subject's brain MRI data, with each voxel's anatomical classification given; the above modeling associates each atlas voxel with a distribution of tissue density variations, and a distribution of geometric variations between individuals. These distributions enrich the atlas into a statistical atlas that embodies the knowledge of anatomical variations in a population. Figure 6 illustrates the concept.



**Figure 6.** A statistical atlas.

## 5. Registration Using the Statistical Atlas

Using the statistical models as prior knowledge, the registration between a subject and the atlas can be formulated as finding the deformation  $D$  that gives the highest posterior probability  $P(D|dI)$ . According to Bayes rule,  $P(D|dI)$  can be expressed as:

$$P(D|dI) = \frac{P(dI|D)P(D)}{P(dI)} \quad (3)$$

Finding the highest  $P(D|dI)$  becomes maximizing the right hand side of equation (3). Here  $P(dI)$  is a constant for two given image volumes, and the numerator has the same maximum as its logarithm. Substituting from equations (1) and (2) and taking logarithms we obtain:

$$\log P(D|dI) = \log \frac{1}{\sqrt{2\pi}\sigma} - \frac{(dI - \mu)^2}{2\sigma^2} + \log \frac{1}{\sqrt{(2\pi)^3|\Phi|}} - \frac{(\bar{\delta} - \bar{\omega})^T \Phi^{-1} (\bar{\delta} - \bar{\omega})}{2}$$

hence maximizing  $P(D|dI)$  is equivalent to minimizing the term:

$$\frac{(dI - \mu)^2}{2\sigma^2} + \frac{(\bar{\delta} - \bar{\omega})^T \Phi^{-1} (\bar{\delta} - \bar{\omega})}{2} \quad (4)$$

We use gradient descent to find the deformation that minimizes (4). The 3-D gradient,  $\nabla$ , at each step of the descent is given by the first order derivative of (4):

$$\nabla = \frac{dI - \mu}{\sigma^2} \nabla I + \Phi^{-1} (\bar{\delta} - \bar{\omega}) \quad (5)$$

where  $\nabla I$  is the 3-D image gradient, which is a function of the voxel's position. Since  $\sigma$  and  $\Sigma$  can have small values, we add a stabilizing factor to  $\sigma^2$  and to the diagonal elements of  $\Phi$  to regularize the gradient. The empirical value we use is  $\|\nabla I\| + 1$ . The 3-D shift  $\delta D$  is then:

$$\delta D = -\lambda \nabla \quad (6)$$

here  $\lambda$  is a step size constant. In this way, each voxel is guided to search for a counterpart so their match is most probable according to the statistics gathered from a population. We apply 3-D Gaussian smoothing to the voxels' 3-D displacements after each iteration to smooth the deformation. This compensates for the fact that the dependence between the deformation of neighboring voxels is not modeled in the statistical atlas.

This algorithm differs from the previously developed hierarchical deformable registration algorithm [13] in the measurement of the goodness of the voxel deformation flow. In this method, we maximize the posterior probability of the current deformation using statistics gathered from a population, whereas in the previous algorithm we minimize the intensity difference

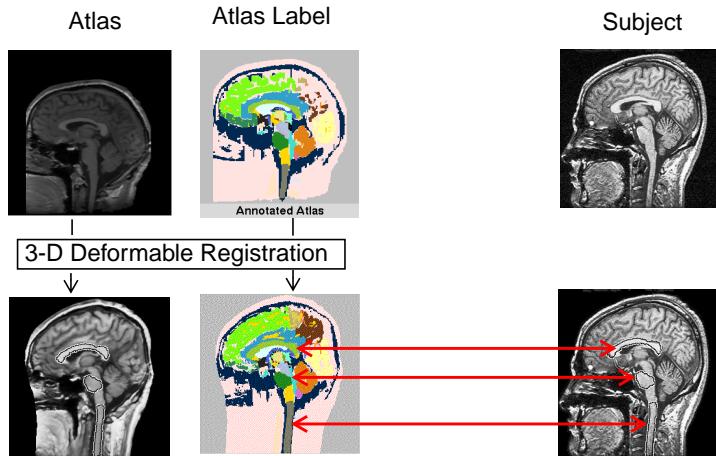
between spatially corresponding voxels in the atlas and the subject volume. Before undergoing deformation, both algorithms globally align the two image volumes to eliminate extrinsic variations<sup>[13]</sup>.

## 6. Performance of Registration Using the Statistical Atlas

We evaluate the effectiveness of our model of anatomical variations by comparing registration using the statistical atlas, and registration using the original atlas.

### 6.1. Evaluation Metric

Since each voxel in the atlas is labelled with the anatomical structure that contains it, when we register the atlas with a subject, we can then assign the label to the corresponding voxel in the subject. This creates a *customized atlas* which contains classifications of the subject's anatomical features. Figure 7 illustrates this process. Given the *ground-truth* classification of the subject's anatomical structures, we can evaluate the quality of the registration by assessing the voxel classification accuracy.



**Figure 7.** Illustration of classifying a subject's anatomical structures through registration with the atlas.

Currently we have 40 subjects' brain MRIs that have expert classification of one structure, the corpus callosum, in one plane, the mid-sagittal plane. Figure 8 shows an example. They are not part of the training set, and are used as the test set.



**Figure 8.** Expert classification of corpus callosum in the mid-sagittal plane of a brain MRI.

Our error metric is the ratio between the number of mislabelled voxels and the number of expert labelled voxels. Mislabelled voxels include those labelled as corpus callosum in the customized atlas but not by the expert, or vice versa.

### 6.2. Registration Using the Intensity Statistics Model

First, we assess the effectiveness of the intensity statistics model alone with constant geometric prior probability. The maximization problem in equation (3) simplifies to  $P(D|dI) \propto P(dI|D)$ , and equation (5) becomes:

$$\nabla = \frac{dI - \mu}{\sigma^2} \nabla I$$

We apply this method to the test set, and compute the ratio of mislabelled voxels for all volumes. We find an overall mislabelled voxel ratio of 3.8%. This is a 14% error reduction over the algorithm with no knowledge guidance<sup>[13]</sup>, which has an

error rate of 4.4%.

### 6.3. Registration Guided by the Geometric Statistics Model

In this experiment, we assess the effectiveness of the geometric statistics model alone, with the intensity statistics assumed to be a constant. The optimum deformation maximizes  $P(D)$ . The 3-D gradient at each step is the second term in equation (5):

$$\nabla = \Phi^{-1}(\vec{\Delta\vartheta} - \vec{\omega})$$

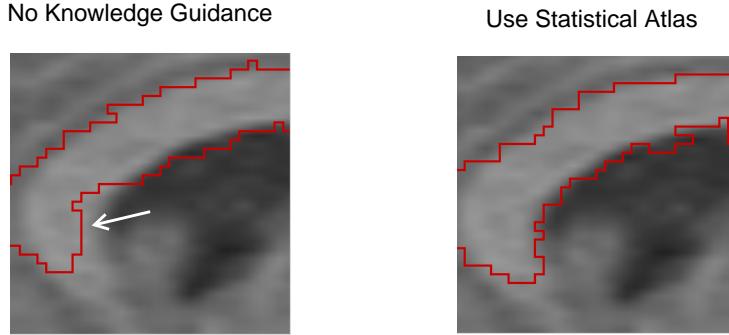
however,  $\nabla$  alone is insufficient to determine the deformation because it ignores the images being registered. To combine the prior model prediction and the image gradient information, we use their inner product to obtain a 3-D deformation gradient  $\tilde{\nabla}$ :

$$\tilde{\nabla} = \left( \frac{\nabla}{\|\nabla\|} \bullet \frac{\nabla I}{\|\nabla I\|} \right) \nabla \quad (7)$$

this balances the influence of the prior distribution and the fidelity to the image data. We apply this method to the test set, and it yields an overall mislabelled voxel ratio of 4.05%, which is an 8% error reduction

### 6.4. Registration Using the Statistical Atlas

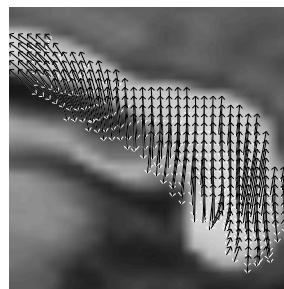
The statistical atlas combines the strength of the intensity and geometric prior, as derived in equation (5). When applied to the test set, it has an overall mislabelled voxel ratio of 3.6%, which is a 23% error reduction over the algorithm with no knowledge guidance. Figure 9 shows an example of improved registration using the statistical atlas



**Figure 9.** Comparison of registration results using the original and the statistical atlas.

## 7. Neighborhood Context

The voxel-based statistics models are efficient at modeling anatomical variations. In reality, however, the deformation of neighboring voxels are not independent. Figure 10 shows an example of the deformation flow overlaid on the image data. Note that the deformation flow is smooth and congruous locally. A more comprehensive model should consider the dependencies between the 3-D deformation of neighboring voxels.



**Figure 10** A close-up on the deformation flow overlaid on the image data. Note the local smoothness of the deformation.

Modeling neighborhood context can be a direct higher dimensional extension of the voxel-based statistics models. Consider a 3-D neighborhood of  $N \times M \times K$  centered at an atlas voxel; the intensity distribution of this neighborhood under a specific

deformation  $D$  is modeled as an  $\mathfrak{N}$  dimensional Gaussian distribution, where  $\mathfrak{N}$  equals NxMxK:

$$P(\vec{d}l|D) = \frac{1}{\sqrt{(2\pi)^{\mathfrak{N}}|\Sigma|}} e^{-\frac{(\vec{d}l - \vec{\mu})^T \Sigma^{-1} (\vec{d}l - \vec{\mu})}{2}} \quad (8)$$

here  $\vec{d}l$  is an  $\mathfrak{N} \times 1$  vector of intensity differences between the corresponding neighborhoods in the subject volume and the atlas.  $\vec{\mu}$  is the  $\mathfrak{N} \times 1$  mean vector of the neighborhood intensity difference distribution;  $\Sigma$  is the  $\mathfrak{N} \times \mathfrak{N}$  covariance matrix of the intensity difference distribution.  $\vec{d}l$  has been adjusted for extrinsic intensity variations

Similarly the geometric variations of a NxMxK neighborhood centered at each atlas voxel can be modeled as a  $3\mathfrak{N}$  dimensional Gaussian distribution of the neighborhood's 3-D deformation:

$$P(D) = \frac{1}{\sqrt{(2\pi)^{3\mathfrak{N}}|\Psi|}} e^{-\frac{(\vec{\Delta}\vartheta - \vec{\omega})^T \Psi^{-1} (\vec{\Delta}\vartheta - \vec{\omega})}{2}} \quad (9)$$

where  $\vec{\omega}$  is the  $3\mathfrak{N} \times 1$  mean vector of the neighborhood's 3-D deformation flow;  $\Psi$  is the  $3\mathfrak{N} \times 3\mathfrak{N}$  covariance matrix of the geometric distribution;  $\vec{\Delta}\vartheta$  is the  $3\mathfrak{N} \times 1$  vectors of the neighborhood's 3-D displacement, and it has been adjusted for extrinsic geometrical variations. Note that the voxel-based statistics models are a special case of the neighborhood statistics models with a  $1 \times 1 \times 1$  neighborhood.

## 8. Registration Using Neighborhood Statistics

We follow the same deduction procedure as in Section 5 to achieve the deformation that maximizes the posterior probability  $P(D|dl)$  for a voxel neighborhood; the 3-D gradient of voxels in the neighborhood is:

$$\nabla = \Sigma^{-1}(\vec{d}l - \vec{\mu}) \nabla I + \Psi^{-1}(\vec{\Delta}\vartheta - \vec{\omega}) \quad (10)$$

Theoretically we can implement this algorithm in the same way as the voxel-based statistics models; however, the  $\mathfrak{N} \times \mathfrak{N}$  intensity covariance matrix  $\Sigma$  has  $(\mathfrak{N}(\mathfrak{N}+1))/2$  distinct entries, and the  $3\mathfrak{N} \times 3\mathfrak{N}$  geometric covariance matrix  $\Psi$  has  $(3\mathfrak{N}(3\mathfrak{N}+1))/2$  distinct entries. Our image volumes typically have more than 8 million voxels. Even if all entries in the covariance matrices can be stored as bytes, the covariance information for each  $2 \times 2 \times 2$  voxel neighborhood will require 336 MByte memory. Together with other memory requirements, the dimensionality of our image volumes makes this approach impractical.

To simplify the problem, we consider only the interaction between immediate neighbors. Instead of storing interactions between immediate neighbors, we compute them on the fly. We approximate the voxel-neighbor interaction using the *goodness* of its neighbors' current match according to their prior distributions. Using a weighted-window matching approach, the *goodness* is weighted by the distance between the voxel and the particular neighbor. Therefore, for a voxel neighborhood  $\mathfrak{N}$ , the 3-D gradient determined by neighborhood statistics models is a direct extension of equation (5):

$$\nabla = \sum_{i,j,k} w_{ijk} \left( \frac{dl - \mu}{\sigma^2} \nabla I \right)_{ijk} + \sum_{i,j,k} w_{ijk} [\Phi^{-1}(\vec{\Delta}\vartheta - \vec{\omega})]_{ijk} \quad ijk \in \mathfrak{N} \quad (11)$$

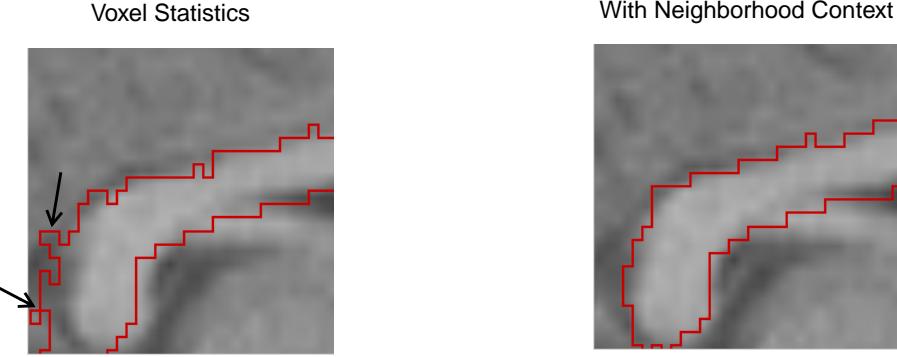
### 8.1. Performance of Neighborhood Statistics Models

We evaluate the effectiveness of neighborhood statistics models in the same fashion as in Section 6. The size of the voxel neighborhood we used is  $3 \times 3 \times 3$ .

**Registration Guided by Neighborhood Intensity Statistics model:** Using only the neighborhood intensity prior model, i.e., the first term in equation (11) achieves an overall error ratio of 3.18%. This is a 27.7% reduction over the algorithm with no knowledge guidance, and a 16% error reduction over registration guided by voxel-based intensity statistics.

**Registration Guided by Neighborhood Geometric Statistics model:** Using only the neighborhood geometric prior model, i.e., the second term in equation (11) gives an overall error ratio of 3.76%. This is a 14.6% reduction over the algorithm with no knowledge guidance, and a 7% error reduction over registration guided by voxel-based geometric statistics.

**Registration Guided by Neighborhood Statistics models:** Using both the neighborhood intensity and geometric prior distributions gives an overall mislabelled voxel ratio of 2.9%. This is a 34% error reduction over the algorithm with no knowledge guidance, and a 20.6% error reduction over registration guided by voxel-based statistics models. Figure 11 shows an example of improved registration using neighborhood statistics models.



**Figure 11** Compare registration results given by voxel-based statistics and neighborhood statistics.

These experiments show that the neighborhood statistics models are significantly more effective than voxel-based statistics models. In the case of intensity statistics model, the neighborhood statistics model nearly doubled the error reduction of the voxel-based statistics model. The geometric statistics model did not seem to be as effective as the intensity statistics model. We attribute this to its higher dimensionality, which would require a larger training set for more accurate model extraction. We expect a complete implementation of the neighborhood statistics guided registration to yield an even greater improvement in performance.

## 9. Related Work

Many researchers have been exploring the characterization of anatomical variations. Bookstein analyzed the underlying statistical and geometric models of morphometrics<sup>[7]</sup>, and applied landmark-based morphometrics to biomedical shape comparison<sup>[8]</sup>. Joshi et al. developed mathematical representations for neuro-anatomically significant substructures of the brain and their variability in a population<sup>[17]</sup>. The “shape” of the substructures are quantified via the construction of templates; the variations are represented by defining probabilistic deformations of the template. Gee et al. formulated the image matching problem under a Bayesian framework, and designed a statistical shape model of the brain extending through the whole organ<sup>[16]</sup>. Their modal basis for describing anatomic variations is derived directly from spatial mappings which bring previous brain samples into alignment with a reference configuration. This is similar to our approach, however, no details on knowledge application has been presented for comparison. Thompson et al. developed a 3-D statistical method for assessing structural variations in a database of anatomic images<sup>[14]</sup>. They created a 3-D probabilistic surface atlas of the human brain, which reflects the observed variability in stereotaxic space of the connected system of anatomic surfaces. Martin et al. characterized neuropathological shape deformations using finite element modeling and statistical analysis<sup>[9]</sup>. Szekely et al. studied flexible parametric shape models by combining the mean contour with a set of eigenmodes of the parameters to characterize shape variations<sup>[6]</sup>. Guimond et al. developed an automatic method to build average intensity and average shape models of the human brain<sup>[10]</sup>. Wang and Staib incorporated statistical contour models into elastic model based 2-D non-rigid registration<sup>[11]</sup>. In their research on segmentation of magnetic resonance imagery, Kapur et al. considered local spatial coherence using a discrete vector valued Markov Random Field model<sup>[12]</sup>. More recent efforts have combined the strength of physically-based models and statistical models for registration and segmentation.

## 10. Conclusion and Future Work

Inter-subject registration is made difficult due to inherent differences between individuals. Characterization of such anatomical variations can help improve registration performance. We extract the patterns of variations in the appearances of brain structures from a training set of 105 T1-weighted MRI. Registration guided by this prior knowledge achieves higher accuracy on a test set of 40 MRI volumes.

We capture the anatomical variations between individuals by registering the training set with a 3-D digital brain atlas, using

a previously developed 3-D hierarchical deformable registration algorithm [13]. This associates each voxel in the atlas with multi-dimensional distributions of anatomical variations in density and geometry. We evaluate statistical properties of these distributions for a neighborhood of each atlas voxel, and embed these statistics into the brain atlas to build a statistical atlas.

Statistical models embedded in the atlas reflect anatomical variations of a population, and thus can function as prior knowledge. It can guide the registration process to tolerate non-pathological variations between individuals while retaining discrimination between different structures. When applied to the test set, the knowledge-guided registration gives an overall voxel misclassification rate of 2.9%; this is a 34% improvement over the performance of the algorithm without knowledge guidance. Experiments have also shown that statistical models that incorporate local spatial congruity are more effective than single-voxel-based statistical models.

The statistical atlas was built upon results from a registration algorithm without knowledge guidance. Imprecisions in the results can affect the rigorousness of the statistical models. To improve model accuracy, we propose to build an initial statistical atlas from a small but accurately registered training set, then bootstrap it into a more reliable model. Besides guiding deformable registration, our computational model of anatomical variations can also facilitate quantitative investigation of anatomical differences between populations, and help detect abnormal variations due to pathology.

## 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 Kate Fissell in the Carnegie Mellon Psychology Department, Dr. Daniel Rio in the National Institute of Health, and Dr. Matcheri Keshavan in the Western Psychiatric Institute and Clinic of the University of Pittsburgh Medical School, for the brain MRI data. We owe our gratitude to our colleagues Marie Elm, John A. Hancock, Daniel Morris, and David LaRose 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] Evans et al., "Warping of Computerized 3D Atlas to Match Brain Image Volumes for Quantitative Neuroanatomical and Functional Analysis. *Proceedings of SPIE Medical Imaging*, vol. 1445, pp. 236-246.
- [3] Vemuri et al., "An Efficient Motion Estimator with Applications to Medical Image Registration", *Medical Image Analysis*.
- [4] Bajcsy and Kovacic, "Multiresolution Elastic Matching", *Computer Vision, Graphics, and Image Processing*, vol. 46, pp 1-21, 1989.
- [5] Jean-Philippe Thirion, "Fast Non-Rigid Matching of 3D Medical Images", INRIA, Technical Report No. 2547, May, 1995.
- [6] Szekely 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.
- [7] Bookstein, "Shape and the information in Medical Images: A decade of the Morphometric Synthesis", *Computer Vision And Image Understanding*, vol. 66, No. 2, pp. 97-118, 1997.
- [8] Bookstein, "Landmark methods for forms without landmarks: morphometrics of group differences in outline shape", *Medical Image Analysis*, vol. 1, No. 3, pp. 225-243.
- [9] Martin et al., "Characterization of Neuropathological Shape Deformations", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, No. 2, 1998.
- [10] Guimond et al., "Automatic Computation of Average Brain Models", *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 631-640, 1998.
- [11] Wang and Staib, "Elastic Model Based Non-rigid Registration Incorporating Statistical Shape Information", *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 1162-1173, 1998.
- [12] Kapur et al., "Enhanced Spatial Priors for Segmentation of Magnetic Resonance Imagery", *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 457-468, 1998.
- [13] Chen et al., "Anomaly Detection through Registration", *Pattern Recognition*, vol 32, pp. 113-128, 1999.
- [14] Thompson et al., "High-Resolution Random Mesh Algorithms for Creating a Probabilistic 3-D Surface Atlas of the Human Brain", *NeuroImage*, vol 3, No. 1, pp. 19-34, February, 1996.
- [15] Miller et al., "Statistical methods in computational anatomy", *Statistical Methods in Medical Research*, 6:267-299, 1997
- [16] Gee and Le Briquer, "An Empirical Model of Brain Shape", *Maximum Entropy and Bayesian Methods*, August 4-8, 1997.
- [17] Joshi et al., "On the Geometry and Shape of Brain Sub-Manifolds", *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 11, No. 8, pp. 1317-1343, 1997.