A DEFORMABLE NEUROANATOMY TEXTBOOK BASED ON VISCOUS FLUID MECHANICS*

Gary E. Christensen[†], Richard D. Rabbitt[‡], and Michael I. Miller[†]
[†]Institute for Biomedical Computing and
Electronic Signals and Systems Research Laboratory
[‡]Department of Mechanical Engineering
Washington University, St. Louis, Missouri 63130

*Invited Paper.

Pulished in Prince and Runolfsson, editors, Proceedings of the 1993 Conference on Information Sciences and Systems, pp. 211-216, Johns Hopkins University, March 24-26, 1993

Abstract

This paper demonstrates a novel image recognition method based on Grenanders Global Shape Models. A mathematical textbook is constructed to represent the typical structure of a human brain. Next, a set of probabilistic transformation based on the theory of viscous fluids is applied to the textbook to generate a rich family of brain images. We describe a method that estimates the maximum a posteriori transformation which deforms the coordinate system of the textbook into the coordinate system of a particular individual. Once this transformation is determined, information about the individual brain can be queried by applying the transformation to the symbolic data contained in the textbook.

1 Introduction

Although the gross structure of the human brain is topologically uniform across a population of individuals, its shape varies from one individual to the next. The variability is evident in the global shape and size of the brain as well as in the local orientation, shape, and size of the constituent structures. A fundamental task in understanding the inter subject variability of biological shapes is the construction of mathematically precise models that describe the variability in terms of the underlying structure.

We use the method of Grenander's global shape models to represent the variability and structure of the brain. In this method, the typical brain is represented in via the construction of a neuroanatomy textbook. Shape variability is accommodated by defining probabilistic transformations that are applied to the textbook. By applying the transformations to the textbook, a rich family of brains may be represented using a single textbook.

The neuroanatomy textbook is a registered set of images generated from a typical healthy brain within the population. It consists of two types of images, measured and labeled. The measured set of images can consist of Magnetic Resonance (MR) images, Computer Axial Tomography (CAT) images, digitized photographs of the cyrosectioned brain, etc. The labeled images contain information about the measured images, consisting of structure names, sizes, locations, functionality, dependence on other structures, etc. The textbook is used by estimating the transformation that registers the measured images of the textbook with measured images from a subject's brain. Once this transformation is known, all the information contained in the labeled images of the textbook can be mapped onto the measured images of the subject.

The complexity of the anatomy which must be represented–ventricles, cortical folds, etc.–require transformations which are of high spatial dimension. To accommodate normal brain variation, we formulate transformations which allow for large deformations while maintaining continuity of the textbook. We achieve this by constraining the set of transformations to those consistent with physical deformations of real materials. Transformations such as those based on the theory of elasticity develop restoring forces which are proportional to the square of the deformed distance. Except for the smallest deformations, such elastic transformations prevent the textbook from being fully deformed into the shape of the subject's brain. The short comings of the elasticity model can be overcome by using a viscous fluid model which allows the restoring forces to relax over time. In this paper we present new results that are deformations of a neuroanatomy textbook based on viscous fluid transformations.

2 Deformable Textbook

We use Grenander's shape models to represent the variability and structure of the brain by defining a textbook (template) and a set of probabilistic transformations that are applied to the textbook. We define the textbook as it was defined in [1] and define the set of transformations to have the properties of viscous fluid mechanics [2].

The textbook, $\mathbf{T}(x) = (T_1(x), \ldots, T_N(x))$, is an N-valued function consisting of a set of registered images based on a single continuous coordinate system or domain, Ω. Each component of the textbook corresponds to a measured image or a labeled image from a single normal brain. The discrete images in the textbook are extended to the continuous domain Ω using bilinear interpolation. Likewise, a study $\mathbf{S}(x) = (S_1(x), \ldots, S_M(x))$ is an M-valued function consisting of a set of measured images taken from a patient. The set of images in the study are a subset of the images in the textbook, i.e., $M < N$. We assume that a study image is a sample from a sample path continuous, quadratic mean, Gaussian process with mean given by the deformed textbook image of the same modality.

We use Bayesian estimation to find the *maximum a posteriori* (MAP) transformation that deforms the textbook into the study. The MAP transformation, $\hat{h}(x) = x - \hat{u}(x)$, is the transformation that maximizes the posterior density, $p(u|S) \propto p(S|u)p(u)$, where u is called the displacement field. The likelihood density, $p(S|u)$, models the study images as Gaussian processes with mean given by the deformed textbook images where

$$p(S|u) = \frac{1}{Z} \; exp\left[-\frac{1}{2\sigma^2}\sum_{k=1}^{M}\int_{\Omega}(T_k(x-u(x)) - S_k(x))^2\,dx\right]. \tag{1}$$

The prior density, $p(u)$, is defined to give large probability to viscous fluid transformations and low probability to all other transformations. As shown in [2], it is possible to determine the posterior maximizer condition without having to specify the exact form of the prior density. This maximizing condition is a partial differential equation (PDE) on the instantaneous velocity, v, of the deformation field, u, and is given by

$$\mu\nabla^2 v + (\mu+\beta)\nabla(\nabla\cdot v) = b. \tag{2}$$

The body force, b, is the gradient of the likelihood with respect to u and is given by

$$b = \frac{1}{\sigma^2}\sum_{k=1}^{M}(T_k(x-u(x)) - S_k(x))\nabla T_k(x-u(x)). \tag{3}$$

We use an Eulerian reference frame to track the kinematics of the deformation. In this reference frame we define a spatially-fixed set of points–one located at each voxel in the study–to observe the deformation of the textbook. As the textbook continuously deforms over time and space, points (particles) in the textbook flow through the observation points. At any instant of time, the transformation $h(x) = x - u(x)$ gives the original position of a particle as it passes through the observation point located at x. Since we are using an Eulerian reference frame, the velocity, v, of a particle is related to its displacement, u, by $v = \frac{\partial u}{\partial t} + v \cdot \nabla u$. The term $v \cdot \nabla u$ accounts for nonlinear trajectories of particles as they pass through the observation points.

Notice that Equation 2 is difficult to solve for $\hat{u}$ directly because it is nonlinear in u. However, Equation 2 is easy to solve for the instantaneous velocity, v, because it is linear in v. Thus, we estimate $\hat{u}$ by discretizing time, solving Equation 2 for v at each discrete time point and for each pixel in the study, and then integrating v over time to get $\hat{u}$.

We illustrate the MAP estimation of the displacement field with a 1-D example. We begin by discretizing the time interval $[0, \infty)$ into small increments $0 = t_0 < t_1 < t_2 < \ldots$. The displacement field at time t_0 is defined to be zero. In 1-D, Equation 2 reduces to

$$\alpha\frac{\partial^2 v(x,t_i)}{\partial x^2} = [T(x - u(x,t_i)) - S(x)]\frac{\mathrm{d}T(x-u(x,t_i))}{\mathrm{d}x}. \tag{4}$$

Jacobi relaxation [3] is used to solve Equation 4 for the instantaneous velocity $v(x, t_0)$ of the particles at each observation points. Note that the solution of Equation 4 using Jacobi relaxation is computationally intensive because it is an iterative algorithm. The displacement field at time t_1 at each observation point is then determined from

$$u(x,t_{i+1}) = u(x,t_i) + (t_{i+1} - t_i)\Big(1 - \frac{\partial u(x,t_i)}{\partial x}\Big)v(x,t_i)\,. \tag{5}$$

This process is repeated at each discrete time point until the energy norm associated with the body force approaches zero (or a preset limit close to zero).

3 Experiments

The first experiment is shown in Figure 1. The textbook consisted of a 128×128 image of a circle and

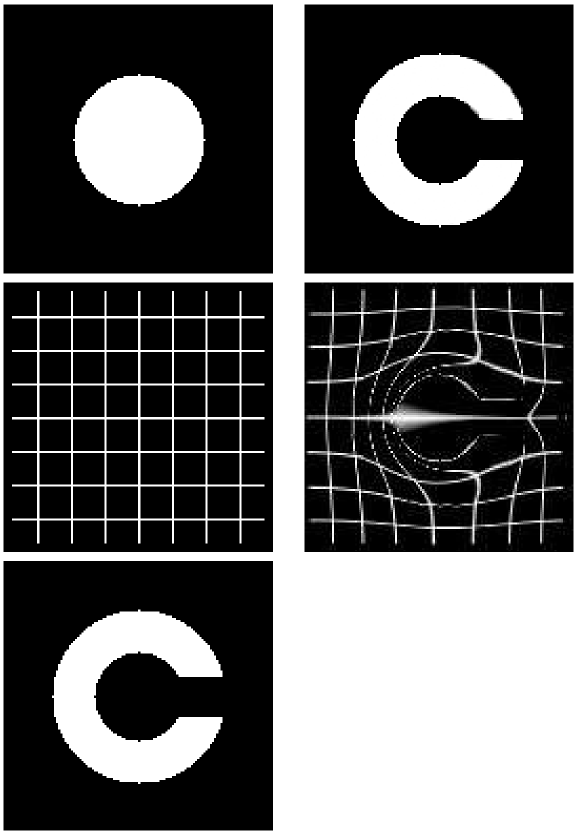

Figure 1: The left column shows the circle textbook which consists of a circle (top) and rectangular grid (middle). The bottom-left panel shows the "C" study. The right column shows the deformed textbook after the circle textbook was deformed into the "C" study.

a grid. The study consisted of an image of a "C" shown in the bottom-left panel of Figure 1. The circle and the "C" were both arbitrarily chosen to be centered. This assumption is not restrictive because as a preprocessing step, the images could have been globally aligned [1]. The radius of the circle is 31 pixels and the inner and outer radius of the "C" is 21 and 41 pixels respectively. The gap of the "C" is 20 pixels wide. The boundary conditions used in this experiment was zero velocity at the four edges of the textbook coordinate system.

Shown in the right column of Figure 1 is the result of deforming the circle textbook into the "C" study. The images of the circle and the "C" were used to find the MAP estimate of the transformation which deforms the circle into the shape of the "C". The right column of Figure 1 shows the result of applying this transformation to the circle (top) and to the grid (middle) of the textbook. There were $4 \times 128 \times 128$ parameters estimated for this transformation, i.e., a displacement and a velocity vector were estimated at each pixel in the study, each has an x and y component. The deformed grid shows how patches of the circle moved and changed shape. One of the properties of the viscous fluid transformation is that it forces the transformation to be a one-to-one mapping of points from the study back to the textbook. We can see the evidence of the one-to-one mapping from the fact that the grid lines in the deformed grid do not cross each other. The reason that the deformed grid lines are not black and white like the original grid is the following. The transformation of the textbook to the study, $h(x) = x - u(x)$, is estimated at each pixel in the study looking back to the continuous image of the grid in the textbook. The continuous image of the grid in the textbook is generated by bilinearly interpolation. A pixel in the deformed grid can have an intensity between black and white if the coordinate of that pixel maps back to a location in the textbook that is between a black and a white pixel.

This example is interesting for a number of reasons. First, in order for the circle to deform into the concave shape of the "C", many points of the circle had to follow nonlinear paths. Deformation procedures that rely on linear approximations such as linear elasticity [1, 4] can not deform the circle into the "C". Another interesting note is the distance that the right side of the circle had to move from its undeformed position to its deformed position. The right side of the circle had to move a distance of 52 pixels (41% of the overall width of the image) to match the interior of the "C". Because the transformation that was estimated was constrained by continuous fluid mechanics, there was no penalty for the distance that a patch of the circle had to move to match a patch of the "C". Elasticity on the other hand enforces a penalty on the transformation that is proportional to the square of the distance that a patch of material moves. Thus, the circle could not have been deformed into the "C" based on elasticity because the right side of the circle would have been held back and prevented from moving the large distance. Finally, notice the sharp corners in the deformed circle. Transformations based on continuous fluid mechanics can have sharp corners. Thus, the circle was able to flow into the sharp corners of the "C". An elastic transformation would have prevented the formation of sharp corners.

The next example shows how we can deform a neuroanatomy textbook into a patient study using viscous fluid transformations. Shown in Figure 2 is the neuroanatomy textbook and the patient study. These images were collected by Dr. Scott Nadel at Duke University. The neuroanatomy textbook consists of a spin-density MR image, a t2-weighted MR image, a grid, and a segmentation of the major nuclei: ventricles, head of the caudate nucleus, thalamus, putamen, other brain matter, and background. The study consists of a spin-density MR image and a t2-weighted MR image. All of these images were 256×256 pixels. It is assumed that the images in the neuroanatomy textbook are in registration with one another and also that the study images are registered with each other. This assumption is valid if the images are collected in registration or if the images were registered with each other in a preprocessing step. The images in this example were chosen because of their similarity in position and orientation. The textbook and study MR images were preprocessed by scaling the intensities from 16-bit integers to 8-bit unsigned integers and the skull was manually removed. In addition, the images in the study were histogram equalized to match the histograms of their respected modalities in the textbook. Histogram equalization is performed so that each structure in both the textbook and the study have roughly the same intensity, i.e., gray matter has roughly the same intensity in both the textbook and the study after histogram equalization.

Notice from Figure 2 that the textbook images and the study images differ globally as well as locally. The textbook brain differs globally from the study in overall size, orientation, and outer contour shape. Locally, the size and shape of constituent structures differs from the textbook to the study. For example, the ventricles of the textbook brain are smaller and of different shape than the ventricles of the study. In addition, the cortical folds of gray matter differ greatly from the textbook to the study.

The MAP transformation that deformed the textbook into the study was estimated in a two step process. First, a global transformation was estimated using the non-rigid, elastic global method described in [1]. Then this global transformation was used as an initial condition to estimate a fluid transformation that matched the local structures of the textbook and the study. Figure 3 shows the result of applying this MAP transformation to the textbook spin-density MR image (top-left), t2-weighted MR image (middle-left), the textbook segmentation (bottom-left), and the textbook grid (middle-right). The result of subtracting the deformed t2 textbook component from the t2 study component and taking the absolute value is shown in the upper-right panel. Figure 3 also contains a hand segmentation of the study (bottom-right) which can be compared to the automatic segmentation (bottom-left). There were $4 \times 256 \times 256$ parameters estimated for the fluid transformation, i.e., a displacement and velocity vector at each pixel in the study.

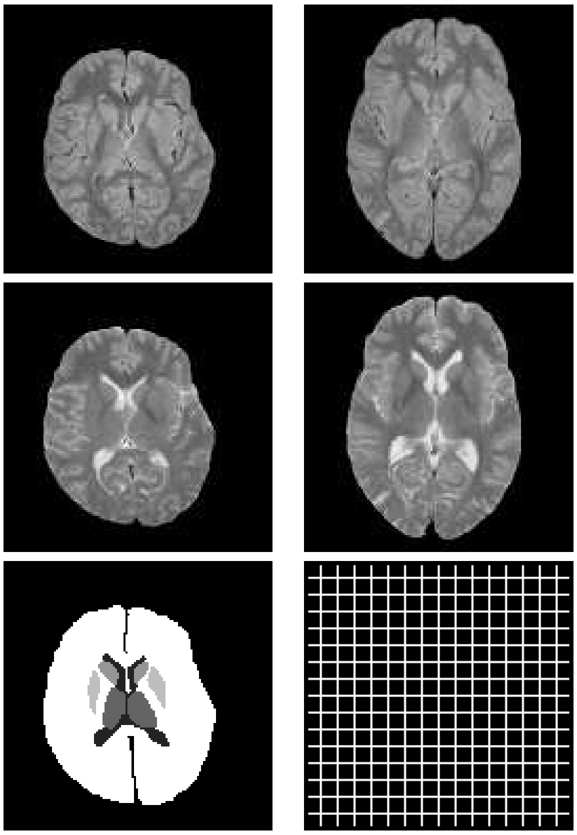

Figure 2: The left column shows three components of a 2-dimensional neuroanatomy textbook. The textbook consists of a Magnetic Resonance (MR) spin-density image (top), a t2-weighted MR image (middle) , a major nuclei segmentation (bottom), and a grid (lower-right). The right column shows a study which consists of a spin-density MR image (top) and a t2-weighted MR image (middle).

We can see from the difference image in Figure 3 that the t2 component of the textbook matched very closely to the t2 component of the study. Ideally the difference image would be black. However, there are a couple of reasons why the difference image need not be black. First, there may be an intensity mismatch between structures in the textbook and study; for example, the ventricles in the study may have a larger intensity than the ventricles in the textbook. Second, both the spin-density and the t2 component of the textbook and study were used to find the transformation. When one modality has more information about a region than another (i.e., a larger mismatch in intensities), the deformation is driven mostly by the modality with the most information. This causes

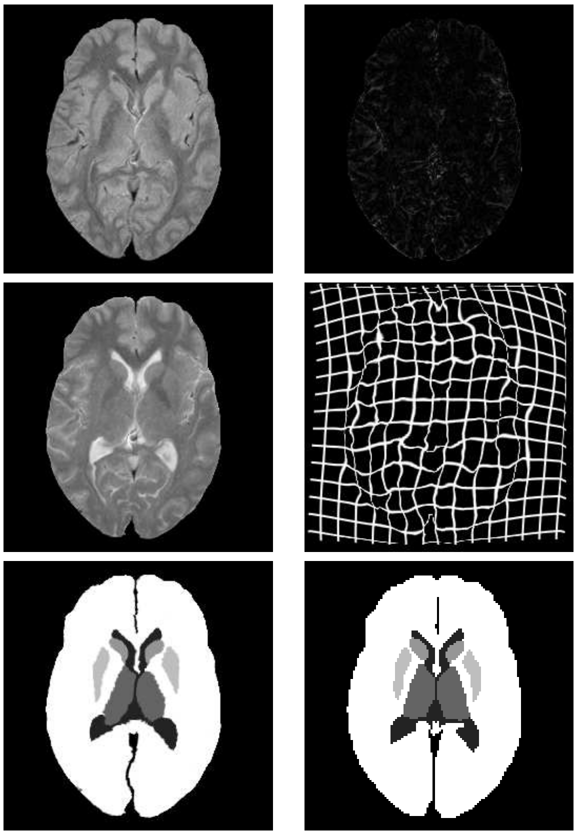

Figure 3: This figure shows the result of applying the estimated transformation to the spin-density (top-left), t2 (middle-left), segmentation (bottom-left), and grid (middle-right) images of the textbook. The top-right panel shows the magnitude difference between the t2 component of the study and the deformed t2 component of the textbook. The bottom-right panel shows a hand segmentation of the study.

an intensity mismatch in the difference image of the modality with less information. A third reason for intensity mismatch is the fluid prior. The prior forces the transformation to be continuous and one-to-one at the expense of creating a mismatch of intensity in a local region. Finally, the presence of different structures in the study and textbook images can cause an intensity mismatch. This can be caused by misalignment in the z direction or it could be caused by an abnormality. The misalignment in the z direction can be corrected by estimating the transformation that deforms the 3-D textbook volume into the 3-D study volume. Intensity mismatch due to abnormal structures is not addressed here.

Figure 3 shows an automatic segmentation of the study. We can see that the automatic segmentation is very reasonable by comparing it to the study images in Figure 2. Notice that the one-to-one mapping property of the fluid transformation kept all of the structures connected and prevented them from being broken apart in the automatic segmentation. The deformed grid also shows that structures were not broken apart because the grid lines are continuous and connected. Also notice the similarity between the automatic segmentation and the hand segmentation of the study. Differences between the automatic and hand segmentations of the study are due to errors in the hand segmentation of the textbook, errors in the hand segmentation of the study, and errors due to partial volume effects. Dann et. al. [5] have pointed out that the fact that there is a large variability in hand segmentations from one medical expert to the next. Using the highest resolution possible for images in the textbook will reduce the errors in hand segmentation of the textbook. Partial volume effects can be minimized by deforming the textbook to the study in 3-D provided the voxel size in the textbook and study is much smaller than the structures of interest.

4 Implementation

The fluid MAP method for estimating the transformation from the textbook to the study was implemented on a 4K DECmpp 12000Sx with MP-2 processors. The DECmpp is a remarketed version of the Maspar computer. It is a massively parallel, single-instruction-multiple-data (SIMD) computer with the processors configured in a 2-D mesh. Each processor element (PE) can communicate with its eight nearest-neighbors via the x-net and can communicate with any PE in the mesh via the global router. In addition, all of the PEs are able to reference data from different locations in their own local memory in parallel.

The fluid deformation algorithm maps well onto DECmpp architecture for a number of reasons. One of the advantages of massively parallel SIMD machines is their ability to compute large numbers of simple calculations very efficiently. The calculation to estimate the velocity and displacement vector at each voxel and each discrete time point is relatively simple, but the large number of vectors that need to be estimated makes this problem very computationally intensive. Secondly, mesh-connected massively parallel computers are very efficient when the calculations only require local interaction of PEs. For the solution of the fluid PDE at each discrete time point, the computation at any PE depends only on its own data and the data contained in its eight nearest- neigh-

Table 1: Execution Times For The Fluid Deformation Algorithm

Experiment	Number of voxels in Study	Time
circle to “C”	128×128	207 sec
2-D neuroanatomy textbook to study	256×256	85 sec
3-D neuroanatomy textbook to study	$256 \times 256 \times 64$	88 min (estimated)

bor PEs. Thirdly, the global router and independent memory addressing capability of the DECmpp allows efficient deformation of the textbook in 2-D or 3-D. In general when the textbook is deformed, PEs need to access data from large distances across the mesh and at different depths in memory. Finally, the DECmpp with MP-2 processors is capable of handling the large amount of computations required to estimate the transformation. At its peak performance, the 4k DECmpp with MP-2 processors is rated at 1.6 GFLOPs for 32-bit operands. In addition, the DECmpp is balanced with respect to communication time and computation time, i.e., data passage between processors takes roughly the same time as one floating point operation. This balance between communication time and computation time is well matched to our algorithm because it also is balanced between communication and computation time.

Table 1 shows the times that were required on the DECmpp to compute the transformations discussed in the two examples. In addition, Table 1 gives an estimated time for computing the 3-D transformation of a $256 \times 256 \times 64$ voxel textbook to a $256 \times 256 \times 64$ voxel patient study using four modalities to force the deformation. The reason the circle example took longer than the 2-D neuroanatomy example is the circle had a larger distance to deform than the neuroanatomy textbook. The larger the distance a textbook must be deformed, the more discrete time steps that are required to deform the textbook.

5 Conclusions

This paper demonstrates a method for representing the variability associated with human neuroanatomies. Two examples were given to demonstrate the process of defining a textbook and estimating the transformation that deforms the textbook into the study. An automatic segmentation of the study brain was generated by applying the MAP transformation estimate to the segmentation in the textbook. In addition, we have shown many of the advantages that fluid transformations have to offer, such as, allowing long distance, nonlinear trajectory deformations, insuring one-to-one deformations, and allowing sharp corners in the deformation if necessary. We have also presented a high speed computer implementation which solves the computationally intensive fluid deformation problem.

6 Acknowledgments

We would like to thank Dr. Scott Nadel from Duke University for supplying the MR images used in this work. G.E. Christensen was supported by NIH DRR-1380. R.D. Rabbitt was supported by the NSF under Presidential Young Investigator Award 8552518. M.I. Miller was supported by NIH-NCRR-RR01380, ONR N00014-92-J-1418, ARO P-29349-MA-SDI, and a grant from the Digital Equipment Corporation.

References

[1] M.I. Miller, G.E. Christensen, Y. Amit, and U. Grenander. A mathematical textbook of deformable neuro-anatomies. *Proceedings of the National Academy of Science*, accepted January 1992.

[2] R. D. Rabbitt, G. Christensen, and M. I. Miller. Deformable 3d templates using fluid dynamics and large deformation kinematics. *To be submitted to the IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1992.

[3] J.C. Strikwerda. *Finite Difference Schemes and Partial Differential Equations*. Wadsworth and Brooks/Cole, 1989.

[4] R. Bajcsy and S. Kovacic. Multiresolution Elastic Matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.

[5] R. Dann, J. Hoford, S. Kovacic, M. Reivich, and R. Bajcsy. Evaluation of Elastic Matching Systems for Anatomic (CT, MR) and Functional (PET) Cerebral Images. *Journal of Computer Assisted Tomography*, 13(4):603–611, July/August 1989.