M.I.T. Media Laboratory Perceptual Computing
Technical Report No. 331 May 1995

Submitted to IEEE PAMI, February 1995

Characterization of Neuropathological Shape Deformations

John Martin, Alex Pentland, Stan Sclaroff, and Ron Kikinis

Surgical Planning Laboratory, Department of Radiology
Brigham and Women's Hospital
75 Francis Street, Boston, MA 02115

Perceptual Computing Section, The Media Laboratory
Massachusetts Institute of Technology
20 Ames Street, Cambridge, MA 02139

Email: martin@bwh.harvard.edu, sandy@media.mit.edu,
sclaroff@cs.bu.edu, kikinis@bwh.harvard.edu

Abstract

We present a framework for analyzing the shape deformation of structures within the human brain. A mathematical model is developed describing the deformation of any brain structure whose shape is affected by both gross and detailed physical processes. Using our technique, the total shape deformation is decomposed into analytic modes of variation obtained from finite element modeling, and statistical modes of variation obtained from sample data.

Our method is general, and can be applied to many problems where the goal is to separate out important from unimportant shape variation across a class of objects. In this paper, we focus on the analysis of diseases that affect the shape of brain structures. Because the shape of these structures is affected not only by pathology but also by overall brain shape, disease discrimination is difficult. By modeling the brain's elastic properties, we are able to compensate for some of the nonpathological modes of shape variation. This allows us to experimentally characterize modes of variation that are indicative of disease processes.

We apply our technique to magnetic resonance images of the brains of individuals with schizophrenia, Alzheimer's disease, and normal-pressure hydrocephalus, as well as to healthy volunteers. Classification results are presented.

1 Introduction

Various neurological disorders affect the gross anatomical shape of different brain structures. These changes have been studied for several decades, using both postmortem and invasive *in vivo* methods. Recent advances in the contrast and resolution of magnetic resonance (MR) scanners now make it possible to study these shape effects *in vivo* and noninvasively, with the potential for better diagnosis and treatment. Our aim is to quantitatively describe these pathological shape deformations.

Previous studies of neuropathological morphology suffer from two drawbacks. First, these studies have used just linear [1-5], planar [6], and/or volumetric [5,7-17] measurements in order to characterize neuropathological shape changes. Research that has used more general shape measures has been qualitative, e.g. having a user manually grade the uniformity of corpus callosum thinning [3][18] and/or smoothness [3]. None of these previous shape descriptions is both general and quantitative.

The second drawback of previous work involves the method of normalizing for nonpathological interpatient differences. These differences are a result of both genetic and environmental factors, which cause biological structures to have a large range of normal variation. To properly study pathological deforma-

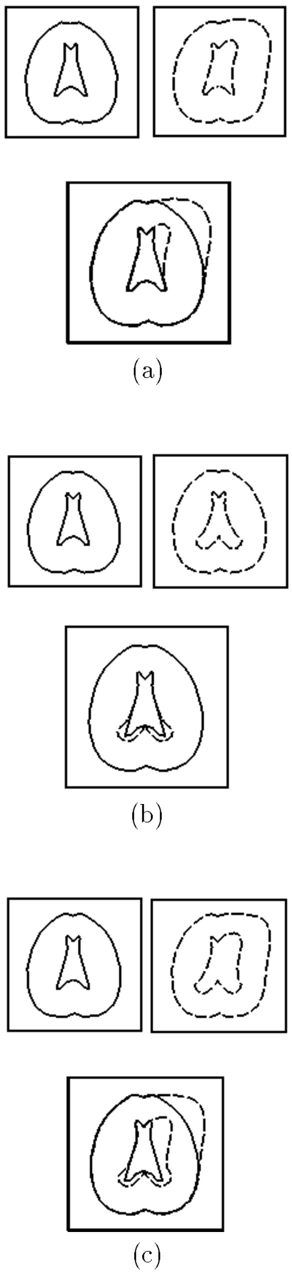

Figure 1: Schematic representation of head and ventricles. (a) Different shaped heads, no ventricular disease present. The only ventricular shape difference is due to the difference in head shape. (b) Same shaped heads, with ventricular disease. The lower tips of the ventricles are expanded due to the disease's physical processes. (c) Different shaped heads, with ventricular disease. The pathological difference in ventricular shape is partially masked out by the nonpathological difference due to head shape.

tions, these nonpathological differences must first be taken into account. Previous studies addressed this by normalizing brain structure measurements for overall brain size. While this is the correct first-order approximation, a more thorough normalization would take into account the complete brain shape.

We overcome these limitations by creating a mathematical framework that (1) separates out disease deformation from deformation due to head shape, (2) uses the complete head shape to normalize cranial contents, and (3) represents pathological deformation in a general and natural manner. Our shape description is in terms of physical and statistical *deformation modes*. These modes can be displayed to show how structures deform due to both head shape and pathology, and can be used in pattern recognition algorithms to classify diseases based on shape changes.

2 The Basic Idea

This paper addresses the general problem of separating out interesting from uninteresting shape deformations in a class of objects. The method we develop has applications in such areas as understanding biological shapes, and in distinguishing important from unimportant shape differences in the encoding of functional classes. In this paper we focus on the specific example of separating out nonpathological shape variation from the pathological deformations caused by various neurological disorders.

Figure 1 demonstrates our framework. Two people with different head shapes will tend to have different ventricular shapes, even in the absence of pathology. This is illustrated in Figure 1a. With disease, however, two people with the same head shape and originally the same ventricular shape will end up with different ventricular shape, as illustrated in Figure 1b. In the most general case, both the effects of head shape and ventricular pathology will be present simultaneously, complicating diagnosis based on ventricular shape. Figure 1c shows this case.

To make this precise, we represent the shape of an average, healthy brain structure as a set of 3-D point positions $\mathbf{X}_A$. This list will contain one entry for every spatial location included in the model. This could be every voxel in the volume, just the surface voxels, or even just a small set of landmarks. The particular choice is an implementation issue. Then, for any particular patient p, the observed deformation $\mathbf{u}_{p,i}$ away from any point $\mathbf{x}_i$ of $\mathbf{X}_A$ can be separated into two distinct components:

$$\mathbf{u}_{p,i}(\mathbf{x}_i) = \mathbf{u}^H_{p,i}(\mathbf{x}_i) + \mathbf{u}^D_{p,i}(\mathbf{x}_i) \quad (1)$$

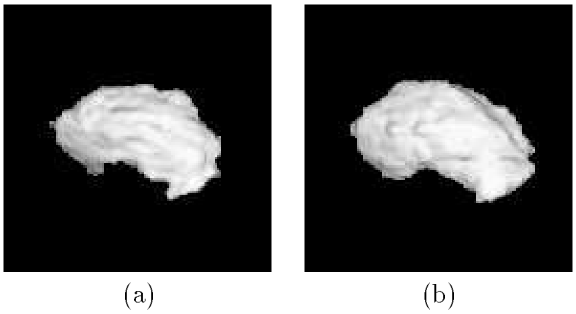

Figure 2: Reconstructions of the left putamen created from MR images. (a) Normal, healthy adult. (b) Patient with schizophrenia.

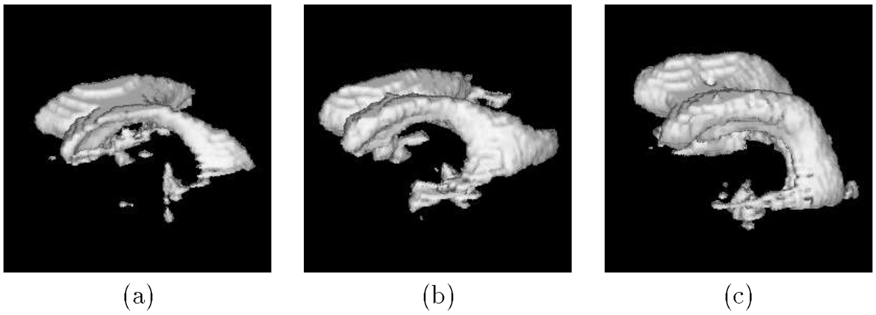

Figure 3: Reconstructions of the cerebral ventricles created from MR images. (a) Normal, healthy adult. (b) Patient with Alzheimer's disease. (c) Patient with normal-pressure hydrocephalus.

where $\mathbf{u}_{p,i}^H(\mathbf{x}_i)$ is the deformation due to global effects that are correlated with overall head shape, and $\mathbf{u}_{p,i}^D(\mathbf{x}_i)$ is the deformation caused by disease and individual local variation. For the entire point set we have

$$\mathbf{U}_p = \mathbf{U}_p^H + \mathbf{U}_p^D \tag{2}$$

where $\mathbf{U}_p$ is the $3V \times 1$ vector

$$\mathbf{U}_p = \begin{bmatrix} \mathbf{u}_{p,1}(\mathbf{x}_1) \\ \mathbf{u}_{p,2}(\mathbf{x}_2) \\ \cdot \\ \cdot \\ \cdot \\ \mathbf{u}_{p,V}(\mathbf{x}_V) \end{bmatrix} \tag{3}$$

and $\mathbf{U}_p^H$ and $\mathbf{U}_p^D$ are defined similarly. V is the number of points in the model.

What is needed then is a method that separates out these two types of deformations, allowing just the pathological deformations $\mathbf{U}_p^D$ to be analyzed. We accomplish this by using the finite element method to create a physical model that describes the macroscopic effects caused by different head shapes. After elastically warping the cranial contents according to this physical model, we are left with residual shape differences across patients that are largely independent of head shape. Once an entire database of patients has been normalized for head shape in this manner, statistical techniques are then used in order to characterize pathological shape variation.

We apply *modal analysis* to the physical modeling, and *principal component analysis* to the experimental observations. Both are eigenanalysis techniques, and represent shape in terms of deformation modes [19][20][21]. These modes represent unique, natural coordinates in which to express the shape and deformation of brain structures.

2.1 An example

To demonstrate our method, we examine deformations in the shape of the left putamen caused by schizophrenia, and deformations in the shape of the cerebral ventricles caused by Alzheimer's disease (AD) and by normal-pressure hydrocephalus (NPH). Figure 2 shows the left putamen of a healthy volunteer and of a patient with schizophrenia, while Figure 3 shows the ventricles of a healthy volunteer, an AD patient, and an NPH patient. Recent studies have shown that schizophrenia can cause the putamen to enlarge [17], and that both AD and NPH cause the ventricles to enlarge [22]. While these studies treated just volumetric changes, we seek to explore other pathological deformations in addition to just volume.

As a reference point for the methodology to be developed in the following sections, we consider here possible ways of classifying patients into the two classes shown in Figure 2. Given a data set consisting of samples from these two categories, the most straightforward classification procedure would be to use just one feature, putamen volume. However, since a person with a larger head will tend to have larger putamen, even if healthy he or she may be misclassified as schizophrenic. Therefore the second procedure to try would be to normalize each person's putamen volume by his or her overall intra-cranial cavity (ICC) volume.

Using each of the above two features, we ran a Gaussian linear classifier on a data set consisting of 13 schizophrenics and 12 normal control subjects. As Table I shows, normalizing for overall ICC volume actually causes the classification rate to slightly decrease. While this decrease in performance is probably mostly due to our small sample size, it also points to possible problems in our normalization. While head *size* is certainly important, the complete head *shape* is really what we ought to use in the normalization.

TABLE I

Putamen Volume Classification Rates

Feature(s)	%Correct
Putamen volume	60
Normalized putamen volume	56

With this in mind, our technique can be viewed as a more sophisticated version of the two features of Table I. Instead of using just brain structure volume, a principal component analysis of the brain structure deformation is calculated, providing other important discriminating features in addition to volume. Also, instead of normalizing for just head size, we normalize for the complete head shape.

2.2 Organization of the paper

The remainder of this paper is organized as follows. Section 3 describes our procedure for head shape normalization. Once a database of patients has been normalized, we then statistically characterize the pathological deformation, as described in Section 4. Section 5 presents the results of applying our technique to two different medical data sets. In Section 6 we discuss our method and experimental results. Section 7 compares our technique to other research, and Section 8 summarizes the work done.

3 Head Shape Normalization

To characterize the global deformations due to head shape, we model the ICC as a linear elastic material, and then set up equations describing its behavior. One reason for using a physically-based model is that we can formulate an approximate physical model for the ICC as a whole. A second and more practical reason is that the finite element implementation of physical modeling provides analytic interpolation functions that allow us to relate deformations at one point to forces and deformations throughout an object. These functions make the task of accurately warping and resampling the data straightforward, allowing us to relate each data set to a standard or normative head shape.

3.1 The finite element method

The most common numerical approach for solving elastic deformation problems is the *finite element method* (FEM) [23]. The major advantage of the FEM is that it uses the Galerkin method of surface interpolation. This provides an analytic characterization of shape and elastic properties over the whole surface, rather than just at the nodes. The ability to integrate material properties over the whole surface alleviates problems caused by irregular sampling of feature points. It also allows variation of the elastic body's properties in order to weigh reliable features more than noisy ones, or to express *a priori* constraints on size, orientation, smoothness, etc. In Galerkin's method, we set up a system of polynomial shape functions that relate the displacement of a single point to the relative displacements of all the other nodes of an object. By using these functions, we can calculate the deformations which spread over the body as a function of its constitutive parameters.

In the *isoparametric* FEM formulation, polynomial shape functions $\mathbf{H}$ are defined in a parametric space $\mathbf{r} = (r, s, t)^T$, with both positions and displacements mapped from parametric to element coordinates using the same shape functions:

$$\mathbf{x}(\mathbf{r}) = \mathbf{H}(\mathbf{r})\hat{\mathbf{X}} \quad (4)$$

$$\mathbf{u}(\mathbf{r}) = \mathbf{H}(\mathbf{r})\hat{\mathbf{U}}. \quad (5)$$

Here $\hat{\mathbf{X}}$ and $\hat{\mathbf{U}}$ denote the nodal position and displacement vectors, respectively, and are defined in the element (object) coordinate system, $\mathbf{x} = (x, y, z)^T$ is any point in the element (object), and $\mathbf{u}$ is the displacement at $\mathbf{x}$. Note that although $\mathbf{u}$ is the displacement in the element coordinate system, because $\mathbf{x}$ is a function of $\mathbf{r}$, $\mathbf{u}$ can be written as a function of either $\mathbf{x}$ or $\mathbf{r}$. (Throughout this paper, a vector with a "hat" ($\hat{}$) denotes a set of FEM nodal positions or displacements, while a vector without a "hat" denotes positions or displacements at a set of non-nodal points.)

For most applications it is necessary to calculate the strain due to deformation. Strain $\boldsymbol{\epsilon}$ is defined as the ratio of displacement to the actual length. The polynomial shape functions can be used to calculate the strains ($\boldsymbol{\epsilon}$) over the body provided the displacements at the node points are known:

$$\boldsymbol{\epsilon}(\mathbf{x}) = \mathbf{B}(\mathbf{x})\hat{\mathbf{U}} \tag{6}$$

where the strain displacement matrix $\mathbf{B}$ is computed by taking the appropriate derivatives of the interpolation matrix $\mathbf{H}$. Because $\mathbf{B}$ is a function of $\mathbf{x}$, and $\mathbf{H}$ is a function of $\mathbf{r}$, the chain rule must be invoked in order to perform the differentiation. This requires the use of the Jacobian matrix $\mathbf{J}$:

$$\mathbf{J} = \begin{bmatrix} \frac{\partial x}{\partial r} & \frac{\partial y}{\partial r} & \frac{\partial z}{\partial r} \\ \frac{\partial x}{\partial s} & \frac{\partial y}{\partial s} & \frac{\partial z}{\partial s} \\ \frac{\partial x}{\partial t} & \frac{\partial y}{\partial t} & \frac{\partial z}{\partial t} \end{bmatrix}. \tag{7}$$

As mentioned earlier, we need to solve the problem of deforming an elastic body subjected to external forces. This requires solving the equilibrium equation

$$\mathbf{K}\hat{\mathbf{U}} = \mathbf{R} \tag{8}$$

for the set of nodal displacements $\hat{\mathbf{U}}$. Here $\mathbf{R}$ is the load vector whose entries are external forces acting on the nodes, and $\mathbf{K}$ is the stiffness matrix. $\mathbf{K}$ is computed directly from the strain displacement matrix by integrating over the object's volume:

$$\mathbf{K} = \int_V \mathbf{B}^T \mathbf{C} \mathbf{B} \; dV, \tag{9}$$

where $\mathbf{C}$ is the *material matrix*, which expresses the material's particular stress-strain law. See Bathe [23] for more details on setting up FEM integrals and equations.

3.2 Modal analysis

It is often more convenient to represent the nodal displacement vector $\hat{\mathbf{U}}$ in the *modal* coordinate system, in which displacements are represented as linear combinations of an object's free vibration modes. These modes provide a unique, natural, and compact coordinate system in which to represent shape and shape change, are computationally efficient to calculate, and have convenient robustness properties with respect to sampling irregularities and measurement noise [19].

To compute the free vibration modes, Equation 8 is diagonalized via an orthogonal transform $\boldsymbol{\Phi}$:

$$\hat{\mathbf{U}} = \boldsymbol{\Phi}\tilde{\mathbf{U}} \tag{10}$$

where $\tilde{\mathbf{U}}$ is a vector of generalized displacements in the new coordinate system. The columns of $\boldsymbol{\Phi}$ are the basis vectors of this new coordinate system.

Substituting Equation 10 into Equation 8 and premultiplying by $\boldsymbol{\Phi}^T$ yields

$$\tilde{\mathbf{K}}\tilde{\mathbf{U}} = \tilde{\mathbf{R}} \tag{11}$$

where $\tilde{\mathbf{K}} = \boldsymbol{\Phi}^T\mathbf{K}\boldsymbol{\Phi}$ and $\tilde{\mathbf{R}} = \boldsymbol{\Phi}^T\mathbf{R}$.

The optimal transformation matrix $\boldsymbol{\Phi}$ is derived from the eigenvalue problem

$$\mathbf{K}\boldsymbol{\phi}_i = \lambda_i\boldsymbol{\phi}_i \tag{12}$$

which, for a discretization with N nodes, has $3N$ solutions $(\lambda_1, \boldsymbol{\phi}_1), (\lambda_2, \boldsymbol{\phi}_2), ..., (\lambda_{3N}, \boldsymbol{\phi}_{3N})$. For dynamic systems, these eigenvectors are called the free vibration modes of the system, with the corresponding eigenvalues giving the square of the vibrational frequency.

We see then that the transformation matrix $\boldsymbol{\Phi}$ has for its columns the eigenvectors of $\mathbf{K}$,

$$\boldsymbol{\Phi} = [\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \boldsymbol{\phi}_3, \cdots, \boldsymbol{\phi}_{3N}], \tag{13}$$

and that $\tilde{\mathbf{K}}$ is a diagonal matrix with the eigenvalues on its diagonal:

$$\tilde{\mathbf{K}} = \boldsymbol{\Phi}^T\mathbf{K}\boldsymbol{\Phi} = \boldsymbol{\Lambda} = \begin{bmatrix} \lambda_1 & & & & \\ & \lambda_2 & & & \\ & & \cdot & & \\ & & & \cdot & \\ & & & & \cdot \\ & & & & & \lambda_{3N} \end{bmatrix}. \tag{14}$$

Because the stiffness matrix has been diagonalized, the resulting system of equations is decoupled and therefore computationally much simpler. Also, the high frequency modes often can and should be discarded for two reasons. First, the low frequency (low eigenvalue) modes contain more information than the high frequency (high eigenvalue) modes in the sense that their amplitudes are larger and therefore for object discrimination they are typically more powerful. Second, because of noise considerations the low frequency modes are more reliably estimated than the high frequency modes.

3.3 FEM model of average ICC

In order to create an FEM model of the average ICC, we first construct a voxel-based average from patient data sets. This is done by automatically segmenting the ICC from each of the data sets [24], performing a rigid body alignment between all the extracted ICCs [25], and then averaging spatial occupancy over the aligned ICCs.

Once the average ICC is constructed, a deformable FEM model is then warped so that its shape approximates that of the average ICC. This procedure starts by extracting the M surface voxels of the average ICC:

$$\mathbf{X}_A^{ICC} = [x_1, y_1, z_1, \cdots, x_M, y_M, z_M]^T . \tag{15}$$

Next we attach virtual springs between each of these voxel coordinates and the closest point on the surface of our finite element model. These virtual springs define forces acting on the deformable model:

$$\mathbf{f}_m = \mathbf{d}(\mathbf{x}_m,\ \text{FEM surface}) \tag{16}$$

where $\mathbf{d}(\cdot)$ is the displacement vector between the given point and the nearest point on the FEM surface. In general, a force $\mathbf{f}_m$ will not act at an FEM node; however we can use the FEM interpolation functions $\mathbf{H}$ to distribute the force to the surrounding nodes in a physically meaningful way. The load vector of Equation 8 is therefore constructed as

$$\mathbf{R}_A = \mathbf{H}^T \mathbf{F}_A . \tag{17}$$

where $\mathbf{F}_A$ is a vector consisting of all the individual forces $\mathbf{f}_m$.

Equation 8 can then be solved for the nodal displacements that give the FEM model a shape approximating that of the average ICC:

$$\hat{\mathbf{U}}_A^{ICC} = \mathbf{K}^{-1} \mathbf{R}_A . \tag{18}$$

These nodal displacements can be added to the original nodal positions to obtain the nodal positions of the average ICC:

$$\hat{\mathbf{X}}_A^{ICC} = \hat{\mathbf{X}} + \hat{\mathbf{U}}_A^{ICC} , \tag{19}$$

where $\hat{\mathbf{X}}$ is the original, undeformed nodal position vector of Equation 4.

We can then transform into the modal coordinate system via Equation 10:

$$\tilde{\mathbf{U}}_A^{ICC} = \mathbf{\Phi}^T \hat{\mathbf{U}}_A^{ICC} \tag{20}$$

where the modal coefficient vector $\tilde{\mathbf{U}}_A^{ICC}$ specifies how much of each deformation mode is contained in the shape of the average ICC. If only the modal coefficients $\tilde{\mathbf{U}}_A^{ICC}$ and not the nodal displacements $\hat{\mathbf{U}}_A^{ICC}$ are required, then we can skip solving for $\hat{\mathbf{U}}_A^{ICC}$, and instead solve directly for $\tilde{\mathbf{U}}_A^{ICC}$, using Equation 11. As mentioned earlier, because the system of equations is decoupled in the modal coordinate system, using Equation 11 directly is much faster.

Figure 4 shows several of the resulting physical deformation modes of the average ICC. In order to obtain additional computational advantages, these modes were computed using the "idealized modes" technique described in [19]. It should be noted that while these idealized modes provide a physical framework and a very useful first approximation to the actual physical modes of the ICC, we cannot claim that these are necessarily the same as the actual modes. The use of more recent implementations of modal fitting [26] will recover modes that are closer to the actual ICC modes.

See [19] for a more detailed description of modal analysis. Software that uses the FEM and modal analysis to recover and describe shapes is available from whitechapel.media.mit.edu in the file /pub/modal.tar.Z.

3.4 FEM fitting to patient ICCs

Just as the deformable FEM model was warped to the shape of the average ICC, the model can also be warped to fit the ICC of any particular patient p. First, the ICC surface points $\mathbf{X}_p^{ICC}$ of the patient are extracted. Virtual springs are then attached between $\mathbf{X}_p^{ICC}$ and the deformable model, generating a force vector $\mathbf{F}_p$. Next these forces are distributed to the FEM nodes, creating the nodal load vector $\mathbf{R}_p$. As before, the nodal displacements can now be found:

$$\hat{\mathbf{U}}_p^{ICC} = \mathbf{K}^{-1} \mathbf{R}_p , \tag{21}$$

and then used to find the nodal positions for patient p's ICC:

$$\hat{\mathbf{X}}_p^{ICC} = \hat{\mathbf{X}} + \hat{\mathbf{U}}_p^{ICC} . \tag{22}$$

Alternatively, we can recover the modal amplitudes $\tilde{\mathbf{U}}_p^{ICC}$ via Equation 11.

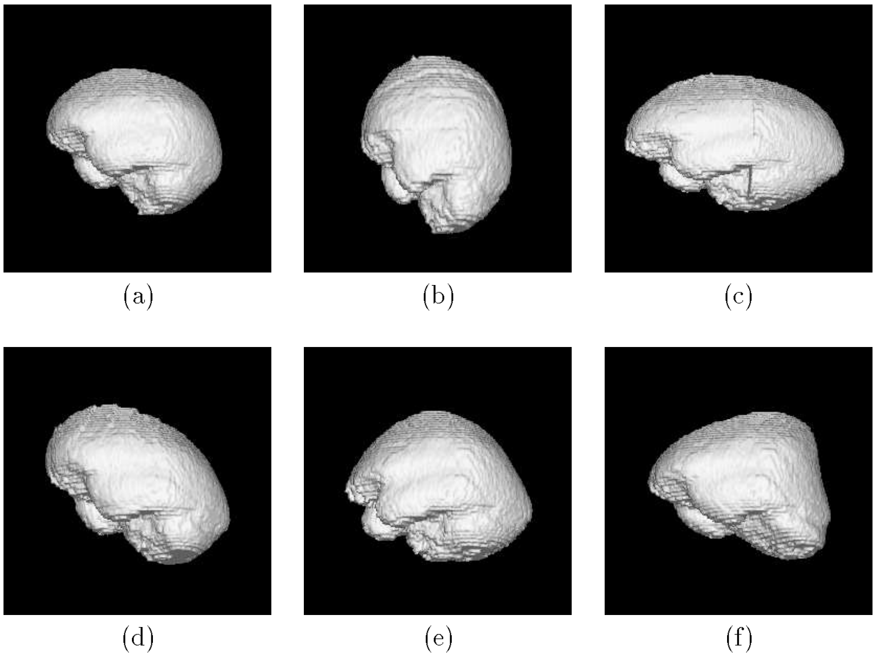

Figure 4: The average ICC (a) and several of the ICC's physical deformation modes (b-f).

Note that although each $\tilde{\mathbf{U}}_p^{ICC}$ (or, equivalently, each $\hat{\mathbf{U}}_p^{ICC}$) represents a different patient p, each patient is warping the *same original set of FEM nodes*. The distinction between patients comes from the particular amount and type of deformation that the set of nodes undergoes; the nodes themselves all start out in the same position. The net effect is that each patient's ICC surface points have now been referred back to the same set of nodes. Thus the original MR sampling differences between patients' ICCs have been removed.

3.5 Warping

The recovered displacement vectors can now be used to normalize each patient's cranial contents in order to account for his or her particular head shape. We have available the set of nodal positions $\hat{\mathbf{X}}_A^{ICC}$ and displacements $\hat{\mathbf{U}}_A^{ICC}$ of the average ICC, as well as $\hat{\mathbf{X}}_p^{ICC}$ and $\hat{\mathbf{U}}_p^{ICC}$ for each patient p. To avoid creating gaps (undefined voxels) when warping, each voxel coordinate in the average ICC coordinate system is mapped into the patient's coordinate system [27]. Because this will produce non-integer coordinates in the patient space, interpolation is necessary in order to calculate an intensity value for the voxel. Repeating this procedure for every voxel position in the average ICC space completely fills up that coordinate system with values from the ICC of patient p.

Recall that the fitting was not done directly from patient ICC to average ICC, but rather from undeformed model to patient ICC, and from undeformed model to average ICC. Thus the required mapping from average ICC space to patient p space must be done in two steps, as described in the next two sections. The final section shows how the other available information - the modal amplitudes $\tilde{\mathbf{U}}_A^{ICC}$ and $\tilde{\mathbf{U}}_p^{ICC}$ - can be used to perform the warping directly in the modal coordinate system.

3.5.1 From average ICC space to parameter space

Given a voxel position $\mathbf{x}$ in the average ICC coordinate system, the first step is to transform it into the parametric coordinate system. This can be done with Equation 4, using the nodal positions of the average ICC:

$$\mathbf{x}(\mathbf{r}) = \mathbf{H}(\mathbf{r})\hat{\mathbf{X}}_A^{ICC}. \tag{23}$$

Note, however, that Equation 23 must be inverted - given $\mathbf{x}$ we need to find $\mathbf{r}$.

This is accomplished as follows. According to Equation 23, once the nodal positions $\hat{\mathbf{X}}_A^{ICC}$ are known, then $\mathbf{x}$ is just a function of $\mathbf{r}$:

$$\mathbf{x} = \mathbf{g}(\mathbf{r}) \tag{24}$$

where $\mathbf{g}$ is the system of three polynomials given in Equation 23. To solve for $\mathbf{r}$, $\mathbf{g}$ must be inverted:

$$\mathbf{r} = \mathbf{g}^{-1}(\mathbf{x}). \tag{25}$$

The solution can be found iteratively using Newton's Method:

$$\mathbf{J}^k\left(\mathbf{x}^{k+1}-\mathbf{x}^k\right)+\mathbf{g}^k=\mathbf{0} \tag{26}$$

where $\mathbf{J}$ is the Jacobian matrix defined in Equation 7, and k represents iteration [28][29].

3.5.2 From parameter space to patient space

Once the parametric coordinates $\mathbf{r}$ are known for $\mathbf{x}$, they must be converted into the coordinate system of patient p. Once again this is accomplished using Equation 4, but this time we use it directly, with the known $\mathbf{r}$ and the ICC nodal positions of patient p:

$$\mathbf{x}'(\mathbf{r})=\mathbf{H}(\mathbf{r})\hat{\mathbf{X}}_p^{ICC}. \tag{27}$$

Once $\mathbf{x}'$ is known, we can simply look into patient p's data set in order to assign a value to $\mathbf{x}$. If the original gray scale data is being warped, then trilinear interpolation can be used to calculate the value. However, because our data sets are segmented, the value is just set to that of the integer coordinates that are nearest to $\mathbf{x}'$. To avoid the aliasing that this introduces, the segmented data can be smoothed before warping.

3.5.3 Modal warping

In our implementation, we recover the modal amplitudes directly, without ever calculating the nodal positions and displacements. As already mentioned, modes offer two important advantages: they decouple the FEM equations to yield improved computational performance, and they provide a unique, canonical coordinate system in which to represent shape.

In the modal coordinate system, instead of Equation 4, the interpolation and warping is done by combining Equations 5 and 10:

$$\mathbf{u}(\mathbf{r})=\mathbf{H}(\mathbf{r})\mathbf{\Phi}\tilde{\mathbf{U}}. \tag{28}$$

To further increase computational efficiency, the polynomial deformations of Equation 28 are approximated by a 3×3 *modal deformation matrix* $\mathbf{D}(\mathbf{r},\tilde{\mathbf{U}})$ [19], which is used to map from parametric to element coordinates:

$$\mathbf{x}(\mathbf{r})=\mathbf{D}(\mathbf{r},\tilde{\mathbf{U}})\mathbf{r}. \tag{29}$$

Thus for a given voxel position $\mathbf{x}$ in the average ICC coordinate system,

$$\mathbf{x}(\mathbf{r})=\mathbf{D}(\mathbf{r},\tilde{\mathbf{U}}_A^{ICC})\mathbf{r} \tag{30}$$

is inverted and solved via Newton's Method to find $\mathbf{r}$, and then

$$\mathbf{x}'(\mathbf{r})=\mathbf{D}(\mathbf{r},\tilde{\mathbf{U}}_p^{ICC})\mathbf{r} \tag{31}$$

is solved directly to find $\mathbf{x}'$.

Equation 30, Equation 31, and the appropriate interpolation scheme can be applied to every voxel position in the average ICC coordinate system. The final result is that locations inside a patient's cranium are displaced to the position that they would occupy if the patient had an average-shaped ICC. By mapping between patient ICC and average ICC in this manner, we account for the geometric differences due to individual head shape, as well as the MR sampling differences between patients.

3.6 Example: Healthy subject with large cranium

Warping of the cranial contents can result in ventricles that are closer to the average. This is demonstrated in Figure 5. Figures 5a and 5b show the average ICC and ventricles, while Figures 5c and 5d show the ICC and ventricles of one of the healthy subjects. This particular subject's ICC is larger than average, particularly in the front-to-back direction (left-to-right in the figure). This ICC shape difference is propagated down to the ventricles, where we see similar shape differences between the two ventricular systems. Calculating the ICC physical deformation modes that warp this subject's ICC to the shape of the average ICC, and then applying that warping to the subject's ventricles, produces the warped ventricles in Figure 5e. As can be seen, these warped ventricles are more similar to the average ventricles in Figure 5b.

4 Characterization of Disease States

4.1 Shape representation

As described in Equations 1-3, our shape representation for a brain structure of a particular patient p is a list of displacements $\mathbf{U}_p$ away from a set of average point positions $\mathbf{X}_A$. These average positions are found by first constructing a volumetric average of the structure from a group of patient data sets, and then extracting the surface voxels.

The displacement vector $\mathbf{U}_p$ is then constructed as follows. Using the original, unwarped data, the surface voxels $\mathbf{X}_p$ of the structure under study are extracted. Then the displacement at a point $\mathbf{x}_i$ of $\mathbf{X}_A$ is computed by finding the nearest point in $\mathbf{X}_p$. By doing this for each of the voxels on the averaged structure's surface, we can compute the displacement vector $\mathbf{U}_p$ of a patient p in the data set.

The above procedure can also be applied to the warped patient data. Because nonpathological deformation is removed by the head shape normalization,

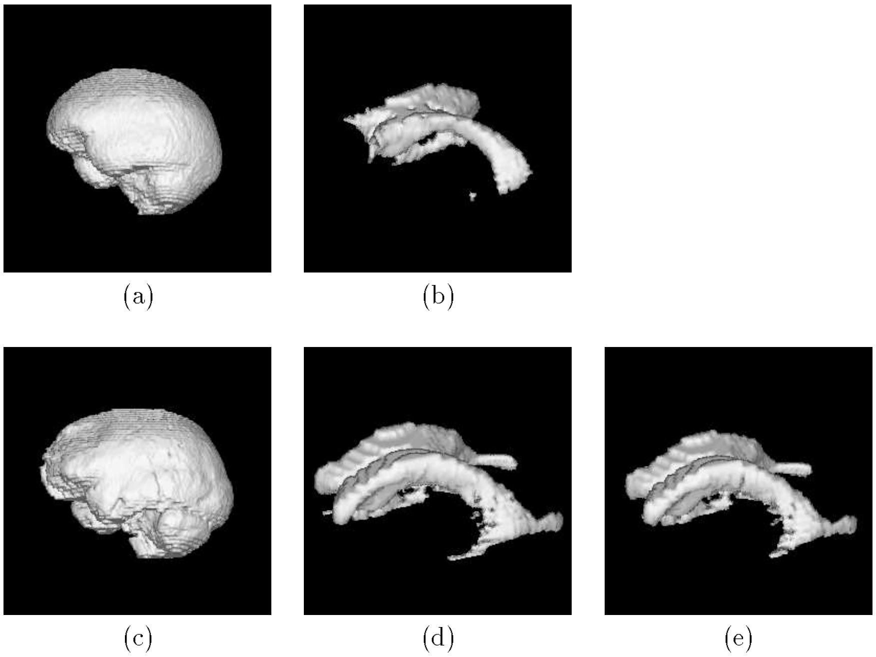

Figure 5: Normalizing ventricular shape by cranium shape. (a) Average ICC (computed from all the data sets). (b) Average normal ventricles (computed from just the healthy subjects). (c),(d) ICC and ventricles of one of the healthy subjects. Compared to (a) this ICC is larger than average, especially in the front-to-back direction (left-to-right from this viewing direction). The ventricles exhibit similar differences as compared to the average ventricles in (b). (e) Subject's ventricles after warping her ICC in (c) to the shape of the average ICC in (a). The ventricles have decreased in size, most notably in the front-to-back direction. Because we have normalized for head shape, the ventricles are now more similar to the average ventricles in (b).

the deformation remaining after warping is due to disease processes (along with any nonpathological differences unaccounted for by head shape). The displacement vector constructed is thus $\mathbf{U}_p^D$ of Equation 2.

We have therefore met our original goal of separating the total deformation into its pathological and nonpathological components. Furthermore, just as with the FEM ICC fitting, because we have referred each patient's particular coordinate system back to the same standard coordinate system, sampling differences between patients' brain structures have been removed. In the remainder of this section, we focus on using the set of pathological displacement vectors $\mathbf{U}_p^D$ to characterize disease.

4.2 Eigenanalysis of shape variation

We can transform the pathological displacement vectors into a coordinate system in which deformations are more naturally represented. This is accomplished through the use of *principal component analysis*.

Principal component analysis is a statistical technique that finds the directions of maximum variability inherent in a data set. When applied to 2-D outline or 3-D surface data, the principal components are called the *eigenshapes* of the structure under study. Unlike the physical modes we have been using throughout this paper, eigenshapes are derived solely from a data set, without the aid of an underlying physical model.

The eigenshapes of our data sets are found as follows. From our P patients, each with a $3V \times 1$ pathological displacement vector $\mathbf{U}_p^D$, the sample covariance matrix is formed as

$$\mathbf{S} = \frac{1}{P-1}\sum_{p=1}^{P}(\mathbf{U}_p^D - \overline{\mathbf{U}}^D)(\mathbf{U}_p^D - \overline{\mathbf{U}}^D)^T, \tag{32}$$

where $\overline{\mathbf{U}}^D$ is the mean of the P displacement vectors. Solving for the eigenvectors of $\mathbf{S}$ yields the principal components, or eigenshapes, of the data set. These eigenvectors are ordered according to their corresponding eigenvalues.

The principal components can be assembled as the columns in a matrix $\mathbf{\Psi}$. Any patient's displacement vector can be written as a linear combination of the principal components:

$$\mathbf{U}_p^D = \mathbf{\Psi}\mathbf{b}_p, \tag{33}$$

where $\mathbf{b}_p$ is the vector of projections onto the principal components for patient p.

4.3 Classifier Design

There are 25 patients in each of our two data sets, and hence 24 principal components. Employing an engineering rule of thumb which states that the top $\frac{1}{4}$ eigenvectors of a sample covariance matrix are reliably estimated, there are six principal components with which to perform classification.

Designing a Gaussian quadratic classifier requires estimating the 6×6 covariance matrix $\mathbf{\Sigma}_i$ of each of the C classes present. However, with only between 7 and 13 members per class, the $\mathbf{\Sigma}_i$ cannot be reliably estimated. The quadratic classifier will consequently be overparametrized, with the result that the training set will not be a good predictor for new cases. We therefore use a Gaussian linear classifier, which assumes that $\mathbf{\Sigma}_1 = \mathbf{\Sigma}_2 = \cdots = \mathbf{\Sigma}_C = \mathbf{\Sigma}$. Thus only the overall $\mathbf{\Sigma}$ has to be estimated, which can be done using all the data sets.

In Gaussian linear classification, a discriminant function is computed for each of the C classes[1]:

$$g_i(\mathbf{b}_p) = 2\mathbf{m}_i^T \mathbf{\Sigma}^{-1}\mathbf{b}_p - \mathbf{m}_i^T \mathbf{\Sigma}^{-1}\mathbf{m}_i \tag{34}$$

where $\mathbf{m}_i$ is the mean of class i and $\mathbf{b}_p$ represents the data set to be classified. In our case $\mathbf{b}_p$ is a vector of the projections of patient p onto the first six principal components (see Equation 33), and $\mathbf{m}_i$ is the 6×1 mean vector computed by averaging the projections over disease class i. The classification rule is then to choose the class i which has the largest g_i, or, equivalently, the maximum probability density when evaluated at $\mathbf{b}_p$ [30].

In order to separate the training stage from the classification stage, we use the *leave-one-out* [31] method. In this procedure, the sample to be classified is withheld from the other samples, which are then used to design the classifier. The held-out sample is then classified. These two steps are repeated for each of the samples, and the results tallied to arrive at the classification rate. Use of this procedure prevents an artificial inflation of the classification rate.

[1] Equation 34 assumes that the prior probabilities of all classes are equal.

In summary, the steps of our classification are:

- For each of the 25 patients:
    1. Use the remaining 24 patients to calculate $\mathbf{\Sigma}$ and the C class means $\mathbf{m}_i$.
    2. Compute the C discriminant functions g_i (Equation 34).
    3. Choose as the winner the class i with the largest g_i.

5 Experimental Results

5.1 Schizophrenia

Thirteen schizophrenic patients and 12 healthy control subjects, matched for gender, age, and handedness, underwent an MR brain scan. As part of an ongoing volumetric study [7][17], the basal ganglia were manually segmented from these scans. Because results of the volumetric study indicated that the basal ganglia of schizophrenics may increase in volume, we decided to examine the basal ganglia for other types of shape changes. We chose the putamen because its relatively large size and simple shape are attractive features when attempting to extract a shape description.

First, using the techniques of Section 3, the cranial contents of each patient were warped in order to normalize the database for head shape. Next, the pathological displacement vector $\mathbf{U}_p^D$ was calculated for each patient's left putamen, as described in Section 4.1. A principal component analysis of the 25 displacement vectors was then performed.

Figure 6 shows the first two putamen principal components. The first mode (Figure 6a-e) is a contrast between the size of the top of the putamen and the size of the bottom, while the second (Figure 6f-j) contrasts the size of the front and back of the putamen (left and right in the figure).

Next we input the top six principal components into our Gaussian linear classifier. Table II shows the results, along with the classification rate using just putamen volume. A 12% improvement in the classification rate occurs when using the putamen principal components instead of just volume.

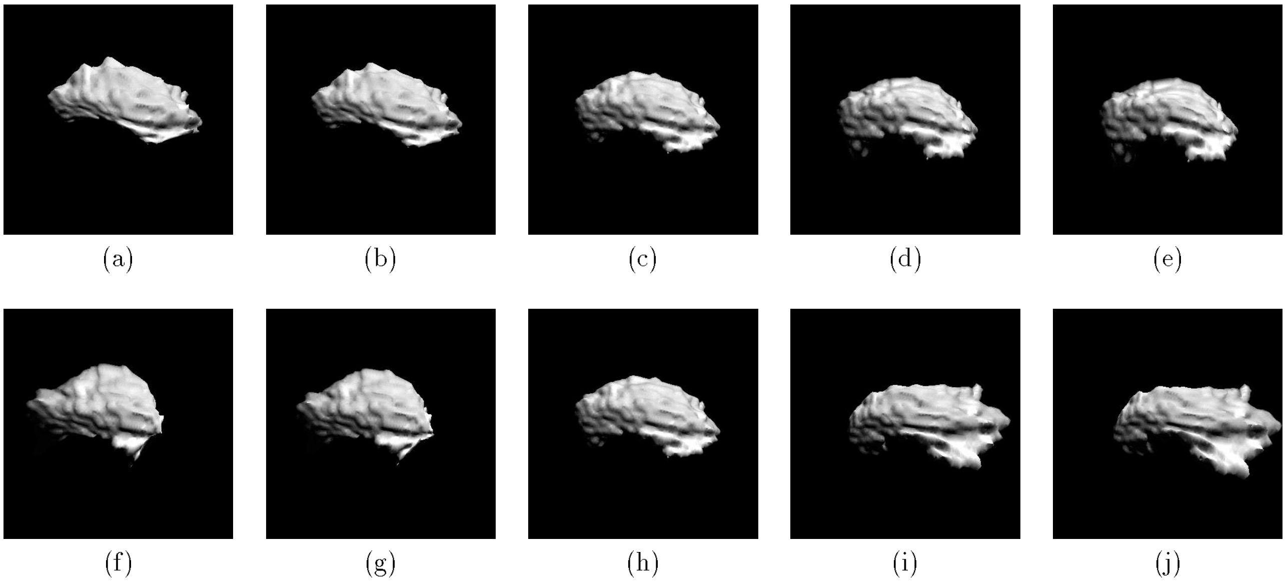

Figure 6: (a)-(e) First principal component of the putamen data set. The amplitude of the mode is increasing from (a) to (e). (f)-(j) Second principal component of the putamen data set. The amplitude of the mode is increasing from (f) to (j).

TABLE II

Putamen Classification Rates

Feature(s)	%Correct
Putamen volume	60
Putamen principal components	72

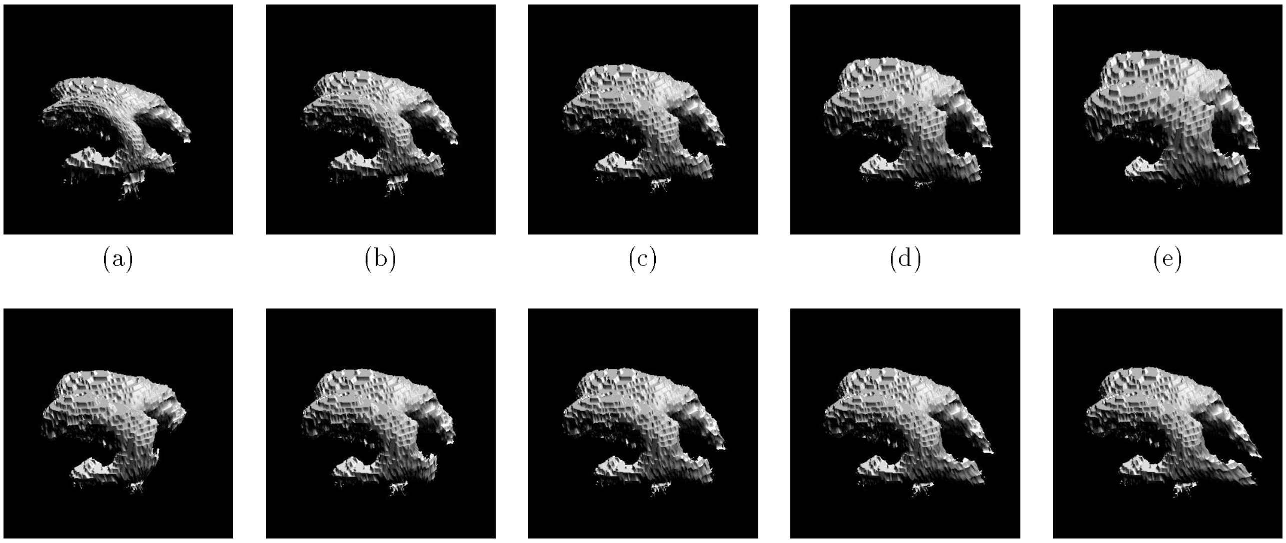

(a) (b) (c) (d) (e)

(f) (g) (h) (i) (j)

Figure 7: (a)-(e) First principal component of the ventricular data set. The amplitude of the mode is increasing from (a) to (e). (f)-(j) Second principal component of the ventricular data set. The amplitude of the mode is increasing from (f) to (j).

5.2 Ventricular disorders

Nine patients with Alzheimer's disease, 7 patients with normal-pressure hydrocephalus, and 9 healthy control subjects, all matched for age, underwent an MR brain scan as part of a previous volumetric study [22] in which the ventricles were segmented using a semi-automatic procedure. Using precisely the same steps that were applied to the putamen, we estimated the principal components of the ventricular data set.

Figure 7 shows two of the principal components. The first mode (Figure 7a-e) is just a measure of the overall size of the ventricles, while the second (Figure 7f-j) is the degree of extension of the posterior horns of the lateral ventricles.

Table III shows the results of running our Gaussian linear classifier on the top six principal components, along with the classification rate obtained using just ventricular volume. As before, principal component analysis has provided a feature space with improved class separability.

TABLE III

Ventricle Classification Rates

Feature(s)	%Correct
Ventricle volume	80
Ventricle principal components	88

6 Discussion

6.1 Effect of head shape normalization

To test the effect of head shape normalization, we computed the putamen principal components *without* first warping all patients' cranial contents to the same model. For the putamen data set, this dropped the classification rate from 72 to 64%, indicating that our normalization did remove some nonpathological putamen deformation.

Since the classification is performed in a six-dimensional space, it is difficult to visualize. We therefore plotted just the top two principal components for all 25 patients in the putamen data set, both without and with the head shape normalization. Figure 8a shows the projections onto the top two modes of the unwarped data, while Figure 8b shows the projections onto the top two modes computed after warping. Note that b_1 and b_2 of Figure 8b are the first two components of the projection vector $\mathbf{b}$ of Equation 33, and that the eigenvectors onto which they are projected were shown earlier in Figure 6. Similarly, a_1 and a_2 of Figure 8a are projections onto the two highest-variance eigenvectors of the sample covariance matrix of the original, unwarped putamen.

Consistent with the decrease in classification rate,

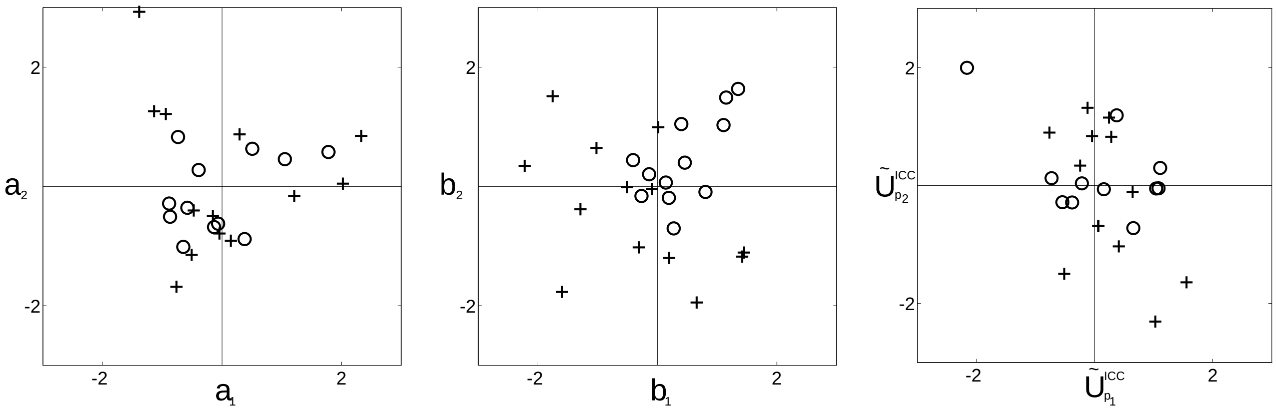

Figure 8: Projections onto modes computed from the schizophrenia database. Each schizophrenic patient is denoted by a $+$, and each healthy volunteer is denoted by a $\circ$. (a) Putamen principal components, computed from the original data. (b) Putamen principal components, computed after first normalizing for head shape. (c) ICC physical deformation modes.

Figure 8a shows less class separability than Figure 8b. This is because the head shape normalization has removed some of the nonpathological shape variation between patients. This nonpathological deformation is represented by the projections onto the first two ICC deformation modes, shown in Figure 8c. As can be seen, the projections onto the ICC modes show little class distinction. This is to be expected, since cranium shape is uncorrelated with disease state.

We repeated this procedure for the ventricular database. This time the classification rate did not drop, staying at 88%. Figure 9 shows the projections onto the top two modes for the unwarped ventricles, the warped ventricles, and the ICC. No improvement is seen in class separability between Figures 9a and 9b.

One possible explanation for the lack of improvement when normalizing for head shape is the already high classification rate (88%) obtainable without removing any nonpathological effects. Coupled with the small sample size, there is little room for improvement. Another factor is our simple nearest-point correspondence scheme. For calculating the correspondence between two ICCs or two putamen, this procedure is adequate. For structures as complicated as the ventricular system, however, nearest-point techniques will provide only a very coarse approximation to the true correspondence. A third possible cause is that the ventricular data set was not controlled for gender. On average the male cranium is larger than the female, but interior structures do not necessarily scale by precisely the same amount [22]. Since normalizing an entire database to one standard head shape does not take into account gender-based shape variation, this gender-based variation may be interfering with the analysis of pathological shape differences.

To further investigate this problem, we constructed a mean square measure of ventricular similarity, by computing the distance between each voxel on the average ventricular surface and the nearest point on a patient's ventricular surface. The sum of these for a particular patient is a measure of the similarity of the patient to the average. If the ICC warping is truly removing some of the variation between ventricles, then this measure should decrease. Figure 10 shows the results for our 25 ventricular patients. Only a moderate decrease is seen with warping.

6.2 An alternative to finite element modeling

In our current implementation, the ICC is modeled as a homogeneous, linear elastic object. This enabled us to use general finite element methods to analytically characterize the entire ICC. Of course the ICC is not homogeneous, and so our simple physical assumptions will lead to inaccuracies. However, it is possible to bypass this dependence on physical assumptions by using the following connection between the eigenshapes calculated using principal component analysis and the physical deformation modes computed via modal analysis.

To derive this relationship, we begin by interpreting Equation 8 and its solution

$$\hat{\mathbf{U}} = \mathbf{K}^{-1}\mathbf{R} \tag{35}$$

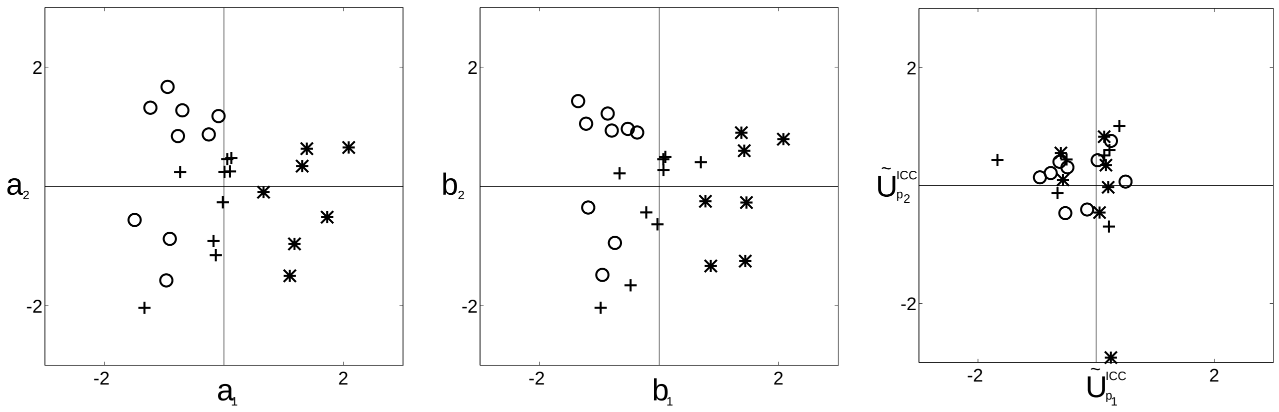

(a) (b) (c)

Figure 9: Projections onto modes computed from the ventricular database. Each Alzheimer's patient is denoted by a $+$, each normal-pressure hydrocephalus patient by an $*$, and each healthy volunteer by a $\circ$. (a) Ventricular principal components, computed from the original data. (b) Ventricular principal components, computed after first normalizing for head shape. (c) ICC physical deformation modes.

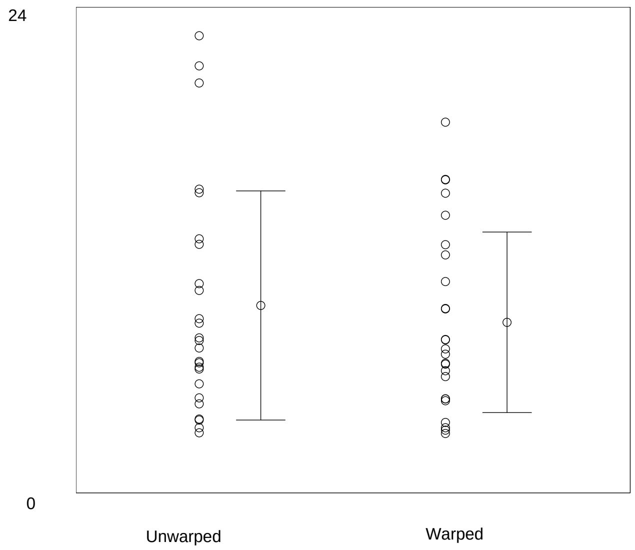

Figure 10: Mean square distance (in mm^2) between ventricles before and after head shape normalization.

from a probabilistic viewpoint. Treating $\hat{\mathbf{U}}$ and $\mathbf{R}$ as random vectors related by the linear transform $\mathbf{K}^{-1}$, we have that

$$\boldsymbol{\Sigma}_{\hat{U}} = \mathbf{K}^{-1}\boldsymbol{\Sigma}_R \left(\mathbf{K}^{-1}\right)^T \tag{36}$$

where $\boldsymbol{\Sigma}_{\hat{U}}$ and $\boldsymbol{\Sigma}_R$ are the covariance matrices of $\hat{\mathbf{U}}$ and $\mathbf{R}$. Because $\mathbf{K}$ is positive semidefinite we can write

$$\boldsymbol{\Sigma}_{\hat{U}} = \mathbf{K}^{-1}\boldsymbol{\Sigma}_R\mathbf{K}^{-1}. \tag{37}$$

Under the assumption that the elements of $\mathbf{R}$ are independent and have variance σ^2, then

$$\boldsymbol{\Sigma}_{\hat{U}} = \sigma^2\mathbf{K}^{-2}. \tag{38}$$

We can form the sample covariance matrix $\mathbf{S} \approx \boldsymbol{\Sigma}_{\hat{U}}$ from a set of observations of $\hat{\mathbf{U}}$, then use Equation 38 to obtain the estimate

$$\mathbf{S} \approx \sigma^2\mathbf{K}^{-2}. \tag{39}$$

Thus by collecting samples of $\hat{\mathbf{U}}$ we can approximate the stiffness matrix $\mathbf{K}$.

This connection leads to several useful observations. First, using a physical model is equivalent to making assumptions about the distribution of samples we expect to see. Not using any model and just collecting data, on the other hand, requires no *a priori* knowledge of this distribution and instead represents an attempt to statistically approximate it through experimental observation.

Second, we have the following result. By applying to $\mathbf{S}$ the orthogonal transform from Section 3.2 that diagonalized $\mathbf{K}$, we have

$$\begin{aligned}
\boldsymbol{\Phi}^T\mathbf{S}\boldsymbol{\Phi} &\approx \sigma^2\boldsymbol{\Phi}^T\mathbf{K}^{-1}\mathbf{K}^{-1}\boldsymbol{\Phi} \\
&= \sigma^2\left(\boldsymbol{\Phi}^T\mathbf{K}^{-1}\boldsymbol{\Phi}\right)\left(\boldsymbol{\Phi}^T\mathbf{K}^{-1}\boldsymbol{\Phi}\right) \\
&= \sigma^2\boldsymbol{\Lambda}^{-1}\boldsymbol{\Lambda}^{-1} \\
&= \sigma^2\begin{bmatrix} \frac{1}{\lambda_1^2} & & & & \\ & \frac{1}{\lambda_2^2} & & & \\ & & \ddots & & \\ & & & & \frac{1}{\lambda_{3N}^2} \end{bmatrix}
\end{aligned} \tag{40}$$

where the second to last step holds because the eigenvalues of the inverse of any nonsingular matrix are just the reciprocals of the eigenvalues of the original matrix. This shows that the orthogonal transform $\mathbf{\Phi}$ also diagonalizes $\mathbf{S}$, which implies that $\mathbf{\Phi}$ consists of the eigenvectors of $\mathbf{S}$. Therefore the eigenvectors of $\mathbf{S}$ converge to those of $\mathbf{K}$, which says that the analytic and estimated modes are the same under the assumption of an independent distribution of loads $\mathbf{R}$. Because of the reciprocal relationship between the eigenvalues of $\mathbf{S}$ and $\mathbf{K}$, the high-variance directions (large eigenvalues) estimated using sample data are the low-frequency directions (small eigenvalues) in a modal decomposition.

Thus we can forgo the reliance on particular physical assumptions and instead compute the stiffness matrix directly from medical imaging data, in the following way. The ICCs from a large database of normal subjects can be used to compute displacement data that relates each patient's ICC surface to a model of an average ICC. Equations 35-39 can then be used to estimate the stiffness matrix. The major drawback to this approach is that without the physical model and the FEM, there will be no physical interpolation functions with which to warp interior structures.

7 Related Work

In the introduction we discussed previous work involving shape measurements of neuropathologies. We then proceeded to present our method of characterizing neuropathological deformation by using both the modes of physical models and of statistical observations.

While our method of using both types of modes is novel, as is its application to neuropathology, both physical modeling and statistical techniques have been used previously in image analysis. In the medical domain, they have been used primarily for registration, segmentation, and/or shape description. Although the goals of these three applications differ, the mathematical techniques employed are often very similar. In this section we review relevant literature from all three application areas, and draw comparisons to our work on shape description.

7.1 Registration

Bajcsy [32] used an elastic model, combining it with cross correlation measures in order to align raw gray scale brain scans with a simplified brain atlas. Collins *et al.* [33][34] also applied a cross correlation measure to align raw patient data with a brain atlas, using both gradient and intensity measures. The allowable deformations were not enforced through physically-based elasticity constraints, but rather by limiting deformations to be on the order of the current scale in their multiresolution approach. Christensen *et al.* [35] implemented both elastic and viscous fluid models of deformation in order to warp patient data to a brain atlas. In their technique, the elastic constraints are used as the prior distribution in a Bayesian formulation. The likelihood function incorporates the agreement between patient data and atlas through a similarity function resembling cross correlation. Minimum mean square estimation (MMSE) is then used to find the posterior distribution, giving the parameters of the elastic transformation from patient to atlas.

Although the above physically-based and probabilistic approaches are formulated quite differently, it has been shown that, under certain conditions, they are in fact equivalent [36][37]. Thus the important differences between the above approaches lie more in the methods of implementation, and thus in speed and convergence properties, than in the particular formulation employed.

While the above approaches have shown promising results, they all suffer from one major drawback. The priors of these models, whether in the form of elastic constants or prior probabilities, are typically chosen in an ad hoc fashion, often for numerical convenience. The result is either that the physical models deform in a physically unrealistic way, or that the prior probabilities in the Bayesian models are poor approximations of the true priors.

7.2 Segmentation

Both physical and probabilistic methods have also been used for medical image segmentation. The primary way in which physical models have been used has been through the use of snakes [38] and their variants. Cohen [39] augmented the original snake formulation with a balloon force to help it avoid local minima. Staib and Duncan [40] used "Fourier snakes", based on a Fourier decomposition of an object's shape. Instead of relying on the elastic constraints of a physical model, they approximated the probability distributions of the Fourier parameters using a training set of manually traced contours, and then applied Bayes rule in order to find the best set of parameter values. Székely *et al.* [41] extended this approach to 3-D using the spherical harmonic technique developed by Brechbühler [42], with statistical deviations in the shape parameters derived using a training set of hand-segmented surfaces.

As in image registration, one of the main problems in segmentation is the construction of realistic prior models. Both Staib and Duncan [40] and Székely *et al.* [41] addressed this by using training examples to form the priors. This method of constructing priors is also an active research area in shape description, as will be described in the next section.

In addition to segmentation, snake models can also be used for shape description. In this domain, however, the mesh-like snake approaches suffer from two problems [19]. First, because the parameters of most snakes can be arbitrarily defined, the recovered shape descriptions are not unique. Second, because the parameters are coupled, the descriptions are also not compact. Both of these drawbacks limit the usefulness of snakes for object recognition.

7.3 Shape Description

7.3.1 Previous Work

Because of the above limitations of mesh-like approaches, researchers have developed other physically-based shape representations. Pentland and Sclaroff [19], on which our work is partially based, represented shape in terms of an object's physical deformation modes. Instead of using the modes of a particular object, Bookstein described shape deformation in terms of the physical deformation modes of an infinite thin plate. Although his original work [43] required corresponding landmarks, more recent efforts [44] have focused on automatically obtaining required points, edges, and surfaces from the image data.

Instead of physically modeling the structure under study, researchers have also sought to obtain shape descriptions directly from sample data. Turk and Pentland [45] have used principle components to describe facial variation and have been able to use this approach to reliably recognize people's faces. Cootes *et al.* [46] used principal components to experimentally describe the modes of variation inherent in a training set of 2-D heart images. Hill *et al.* [47] extended this technique to 3-D and analyzed the cerebral ventricles for purposes of segmentation and representation.

Along with ourselves [48], other researchers have also begun investigating the relationship between physical and experimental modeling. Cootes [21] has examined both physical and statistical shape models, with the goal of smoothly transitioning from a physical to a statistical shape description as more and more data become available. Zhu and Yuille [49] have also considered both physical and statistical shape models, in the context of representing and recognizing objects from their 2-D silhouettes.

7.3.2 Comparison to Our Method

Like modal analysis, all of the above shape descriptors satisfy the requirements of being unique and compact. As argued by Sclaroff and Pentland [20], however, physical deformation modes offer the additional advantage that they are an orthogonal basis for a finite element model. Thus there is a connection with the underlying physics, which is useful for simulation, regularization, and for including *a priori* information about the material properties of the object under study. While Bookstein's thin-plate modes are also physical, they come from a 2-D model and are derived from a finite difference formulation. A finite element model is more general, and has better convergence properties. Also, as mentioned earlier, using the FEM provides interpolation functions that can be used to refer all patients' data back to the same model, thereby removing sampling differences between patients.

A more novel aspect of this paper is the connection between FEM modes and principal component analysis, presented in Section 6.2. Using that development, prior models can be constructed either by making physical assumptions, or through the use of training data. These ideas are similar to the recent work of Cootes [21].

The final original feature of our work is the decomposition of shape deformation into two distinct components. This allows us to remove just the nonpathological deformation, and therefore to focus on pathological morphology. This contrasts with the above work, where only healthy brains are used, and the goal is usually to remove *all* morphological differences between the individual brains. As discussed in the introduction, previous work that has examined pathological brain morphology has used just linear, planar, and/or volumetric shape measurements.

8 Summary

We have presented a new method that addresses the general problem of separating out normal shape variations across a class of objects from those variations that carry special importance. Using both physical modeling and statistical techniques, our method describes shapes in terms of modes of deformation.

When applied to the human brain, our technique is able to separate out pathological from normal shape deformation, allowing better representation and analysis of the deformations due to disease. The representation is in the form of a disease's deformation modes, which provide a very natural basis set in which to examine pathological shape deformation. The analysis suggests that by first discounting the experimentally-derived modes of a brain structure by the physical modes of the intra-cranial cavity, it may be possible to improve disease classification.

Our method was applied to schizophrenia, Alzheimer's disease, and normal-pressure hydrocephalus. The putamen of schizophrenics, although initially very similar to those of normal controls, were easier to differentiate from the control putamen once head shape was taken into account. Conversely, the ventricles of Alzheimer's patients, normal-pressure hydrocephalus patients, and normal controls were somewhat differentiable to begin with, but this separability did not markedly improve when cranial contents were normalized for head shape.

The limitations of our method involve the accuracy of physical models of brain stiffness, the ability to determine the correct correspondence between points on structures, and the degree to which nonpathological and possibly pathological morphology are correlated with overall head shape. Overcoming these limitations will require better implementation, further investigation of the brain's material properties, and shape correlation studies involving large numbers of patients.

In summary, there are two main contributions of this work. First, we have developed a method of shape analysis that is useful for separating out interesting from uninteresting shape variation. By applying modal analysis to the physical modeling, and principle component analysis to the experimental observations, all shape variations were consistently described in terms of deformation modes. Second, when our method is applied to neuropathology, it may be possible to improve disease classification by first normalizing for the physical modes associated with head shape. In addition to serving as features for classification, putamen and ventricle eigenshapes were also displayed in order to illustrate the pathological deformation modes caused by schizophrenia, Alzheimer's disease, and normal-pressure hydrocephalus.

Acknowledgments

The authors would like to thank Irfan Essa for his many helpful comments, Mike Matsumae for providing the ventricular MR brain data, and Martha Shenton for providing the schizophrenic MR brain data.

References

[1] D. Bartelt, C.E. Jordan, E. Strecker, and A.E. James. Comparison of ventricular enlargement and radiopharmaceutical retention: A cisternographic-pneumoencephalographic comparison. *Radiology*, 116:111–115, July 1975.

[2] C.P. Hughes and M. Gado. Computed tomography and aging of the brain. *Radiology*, 139:391–396, May 1981.

[3] T. El Gammal, M.B. Allen, Jr., B.S. Brooks, and E.K. Mark. MR evaluation of hydrocephalus. *American Journal of Neuroradiology*, 8:591–597, July/August 1987.

[4] C. Wikkelsö, H. Andersson, C. Blomstrand, M. Matousek, and P. Svendsen. Computed tomography of the brain in the diagnosis of and prognosis in normal pressure hydrocephalus. *Neuroradiology*, 31:160–165, 1989.

[5] M.J. de Leon, A.E. George, B. Reisberg, S.H. Ferris, A. Kluger, L.A. Stylopoulos, J.D. Miller, M.E. La Regina, C. Chen, and J. Cohen. Alzheimer's disease: Longitudinal CT studies of ventricular change. *American Journal of Neuroradiology*, 10:371–376, March/April 1989.

[6] T. Sandor, M. Albert, J. Stafford, and S. Harpley. Use of computerized CT analysis to discriminate between Alzheimer patients and normal control subjects. *American Journal of Neuroradiology*, 9:1181–1187, November/December 1988.

[7] M.E. Shenton, R. Kikinis, F.A. Jolesz, S.D. Pollak, M. LeMay, C.G. Wible, H. Hokama, J. Martin, D. Metcalf, M. Coleman, and R.W. McCarley. Abnormalities of the left temporal lobe and thought disorder in schizophrenia: A quantitative magnetic resonance imaging study. *The New England Journal of Medicine*, 327(9):604–612, August 1992.

[8] F. Cendes, F. Andermann, P. Gloor, A. Evans, M. Jones-Gotman, C. Watson, D. Melanson, A. Olivier, T. Peters, I. Lopes-Cendes, and G. Leroux. MRI volumetric measurement of amygdala and hippocampus in temporal lobe epilepsy. *Neurology*, 43:719–725, April 1993.

[9] B. Peterson, M.A. Riddle, D.J. Cohen, L.D. Katz, J.C. Smith, M.T. Hardin, and J.F. Leckman. Reduced basal ganglia volumes in Tourette's syndrome using three-dimensional reconstruction techniques from magnetic resonance images. *Neurology*, 43:941–949, May 1993.

[10] H.S. Singer, A.L. Reiss, J.E. Brown, E.H. Aylward, B. Shih, E. Chee, E.L. Harris, M.J. Reader,

G.A. Chase, R.N. Bryan, and M.B. Denckla. Volumetric MRI changes in basal ganglia of children with Tourette's syndrome. *Neurology*, 43:950–956, May 1993.

[11] F. Cendes, F. Andermann, F. Dubeau, P. Gloor, A. Evans, M. Jones-Gotman, A. Olivier, E. Andermann, Y. Robitaille, I. Lopes-Cendes, T. Peters, and D. Melanson. Early childhood prolonged febrile convulsions, atrophy and sclerosis of mesial structures, and temporal lobe epilepsy: An MRI volumetric study. *Neurology*, 43:1083–1087, June 1993.

[12] E.H. Aylward, J.D. Henderer, J.C. McArthur, P.D. Brettschneider, G.J. Harris, P.E. Barta, and G.D. Pearlson. Reduced basal ganglia volumes in HIV-1-associated dementia: Results from quantitative neuroimaging. *Neurology*, 43:2099–2104, October 1993.

[13] S.S. Spencer, G. McCarthy, and D.D. Spencer. Diagnosis of medial temporal lobe seizure onset: Relative specificity and sensitivity of quantitative MRI. *Neurology*, 43:2117–2124, October 1993.

[14] G.D. Cascino, C.R. Jack, Jr., F.W. Sharbrough, P.J. Kelly, and W.R. Marsh. MRI assessments of hippocampal pathology in extratemporal lesional epilepsy. *Neurology*, 43:2380–2382, November 1993.

[15] A.M. Murro, Y.D. Park, D.W. King, B.B. Gallagher, J.R. Smith, F. Yaghmai, V. Toro, R.E. Figueroa, D.W. Loring, and W. Littleton. Seizure localization in temporal lobe epilepsy: A comparison of scalp-sphenoidal EEG and volumetric MRI. *Neurology*, 43:2531–2533, December 1993.

[16] H.S. Soininen, K. Partanen, A. Pitkänen, P. Vainio, T. Hänninen, M. Hallikainen, K. Koivisto, and P.J. Riekkinen, Sr. Volumetric MRI analysis of the amygdala and the hippocampus in subjects with age-associated memory impairment: Correlation to visual and verbal memory. *Neurology*, 44:1660–1668, September 1994.

[17] H. Hokama, M.E. Shenton, P.G. Nestor, R. Kikinis, J.J. Levitt, C.G. Wible, B.F. O'Donnell, D. Metcalf, F.A. Jolesz, and R.W. McCarley. Caudate, putamen, and globus pallidus volume in schizophrenia: A quantitative MRI study. Submitted to *Psychiatric Research: Imaging*, 1994.

[18] C.R. Jack, Jr., B. Mokri, E.R. Laws, Jr., O.W. Houser, H.L. Baker, Jr., and R.C. Petersen. MR findings in normal-pressure hydrocephalus: Significance and comparison with other forms of dementia. *Journal of Computer Assisted Tomography*, 11(6):923–931, November/December 1987.

[19] A.P. Pentland and S. Sclaroff. Closed-form solutions for physically based shape modeling and recognition. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(7):715–729, July 1991.

[20] S. Sclaroff and A.P. Pentland. Modal matching for correspondence and recognition. *IEEE Trans. Pattern Anal. Machine Intell.* To appear.

[21] T.F. Cootes. Combining point distribution models with shape models based on finite element analysis. In *Proc. British Machine Vision Conference*, 1994.

[22] M. Matsumae, A.V. Lorenzo, R. Kikinis, F.A. Jolesz, and P.McL. Black. Intra and extraventricular CSF volumes in normal and ventriculomegalic patients assessed by MR computerized image processing segmentation. In *American Association of Neurological Surgeons*, San Francisco, 1992.

[23] K. Bathe. *Finite Element Procedures in Engineering Analysis.* Prentice-Hall, Inc., 1982.

[24] H.E. Cline, W.E. Lorensen, R. Kikinis, and F.A. Jolesz. 3-D segmentation of MR images of the head using probability and connectivity. *Journal of Computer Assisted Tomography*, 14(6):1037–1045, 1990.

[25] G. Ettinger, E. Grimson, and T. Lozano-Perez. Automatic registration for multiple sclerosis change detection. In *CVPR Workshop on Biomedical Image Analysis*, June 1994.

[26] S. Sclaroff and A.P. Pentland. On modal modeling for medical images: Underconstrained shape description and data compression. In *CVPR Workshop on Biomedical Image Analysis*, June 1994.

[27] G. Wolberg. *Digital Image Warping.* IEEE Computer Society Press, 1990.

[28] G. Strang. *Introduction to Applied Mathematics.* Wellesley-Cambridge Press, 1986.

[29] I.A. Essa, S. Sclaroff, and A.P. Pentland. *Physically-based Modeling for Graphics and Vision.* Chapter in *Directions in Geometric Computing 1992*, Information Geometers, 1992. R. Martin, Ed.

[30] C.W. Therrien. *Decision, Estimation, and Classification: An Introduction to Pattern Recognition and Related Topics*. John Wiley & Sons, 1989.

[31] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, 1990.

[32] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Comput. Vision Graphics Image Process.*, 46:1–21, 1989.

[33] D.L. Collins, P. Neelin, T.M. Peters, and A.C. Evans. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *Journal of Computer Assisted Tomography*, 18(2):192–205, 1994.

[34] D.L. Collins, T.M. Peters, and A.C. Evans. An automated 3D non-linear image deformation procedure for determination of gross morphometric variability in human brain. In *SPIE: Visualization in Biomedical Computing*, pages 180–190, 1994.

[35] C.E. Christensen, R.D. Rabbitt, and M.I. Miller. 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39:609–618, 1994.

[36] R. Szeliski. *Bayesian Modeling of Uncertainty in Low-level Vision*. Kluwer Academic Publishers, 1989.

[37] T.E. Boult, S.D. Fenster, and T. O'Donnell. Physics in a fantasy world vs robust statistical estimation. In *Proc. NSF Workshop on 3D Object Recognition*, New York, NY, November 1994.

[38] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models and 3D object reconstruction. *Int. J. Comput. Vision*, 1:211–221, 1987.

[39] L. Cohen. On active contour models and balloons. *Computer Vision, Graphics and Image Processing: Image Understanding*, 53(2):211–218, 1991.

[40] L.H. Staib and J.S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(11):1061–1075, November 1992.

[41] G. Székely, A. Kelemen, C. Brechbühler, and G. Gerig. Segmentation of 3D objects from MRI volume data using constrained elastic deformations of flexible fourier surface models. Submitted to *CVRMed*, 1995.

[42] C. Brechbühler, G. Gerig, and O. Kübler. Parametrization of closed surfaces for 3-D shape description. *Computer Vision, Graphics and Image Processing: Image Understanding,* to appear, 1995.

[43] F.L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Machine Intell.*, 11(6):567–585, June 1989.

[44] F.L. Bookstein and W.D.K. Green. The biometrics of landmarks and edgels: A new geometry of prior knowledge for medical image understanding. In *Proc. AAAI Symposium on Applications of Computer Vision in Medical Image Processing*, pages 134–137, March 1994.

[45] M. Turk and A.P. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.

[46] T.F. Cootes, D.H. Cooper, C.J. Taylor, and J. Graham. Trainable method of parametric shape description. *Image and Vision Computing*, 10(5):289–294, June 1992.

[47] A. Hill, T.F. Cootes, and C.J. Taylor. A generic system for image interpretation using flexible templates. In *Proc. British Machine Vision Conference*, pages 276–285, 1992.

[48] J. Martin, A.P. Pentland, and R. Kikinis. Shape analysis of brain structures using physical and experimental modes. In *Proc. CVPR*, pages 752–755, 1994.

[49] S.C. Zhu and A.L. Yuille. Forms: A flexible object recognition and modelling system. Technical Report 94-1, Harvard Robotics Laboratory, 1993.