

Characterization of Neuropathological Shape Deformations

John Martin, Alex Pentland, *Member, IEEE Computer Society*,
Stan Sclaroff, *Member, IEEE Computer Society*, and Ron Kikinis

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.

Index Terms—Medical image analysis, shape description, deformable models, finite element method, modal analysis, principal component analysis, eigenanalysis, clustering.

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], [2], [3], [4], [5], planar [6], and/or volumetric [5], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [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 deformations, these nonpathological differences must first be taken into account. Previous studies addressed this by normalizing brain structure measurements for overall brain size.

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. While our current work emphasizes results in the medical domain, the general approach can be applied to a whole cadre of shape categorization and classification problems in which we are given a priori knowledge of shape variation over a particular object class. For instance, the formulation would have fruitful application in tracking and recognition of gesture, facial expression, and gait.

- J. Martin is with Millennium Pharmaceuticals, Inc., 238 Main St., Cambridge, MA 02142. E-mail: martin@mpi.com.
- A. Pentland is with the Media Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: sandy@media.mit.edu.
- S. Sclaroff is with the Computer Science Department, Boston University, 111 Cummington St., Boston, MA 02215. E-mail: sclaroff@cs.bu.edu.
- R. Kikinis is with the Surgical Planning Laboratory, Department of Radiology, Brigham and Women's Hospital, 75 Francis St., Boston, MA 02115. E-mail: kikinis@bwh.harvard.edu.

Manuscript received 10 Feb. 1995; revised 17 Oct. 1997. Recommended for acceptance by S. Dunn.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 106073.

Furthermore, such models can provide a parameterized estimate of principal deformations due to a specific process: manipulation, locomotion, growth, manufacture, disease, wind, heat, etc.

Nonrigid deformation can provide important cues to how shapes are related; however, not all deformations are valid or important in all contexts. Thus, our deformable shape model explicitly includes knowledge of the statistically likely and unlikely deformations. This makes it possible to exploit available *context-specific domain knowledge* in order to reduce the complexity of the nonrigid shape recovery and recognition problem, to gain greater robustness to noise, and to achieve better recognition rates.

2.1 Neurological Shape Deformations

In this paper, we focus on the specific example of separating out nonpathological shape variation from the pathological deformations caused by various neurological disorders. Fig. 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 Fig. 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 Fig. 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. Fig. 1c shows this case.

To make this precise, we represent the shape of an average, healthy brain structure as a set of 3D 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:

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

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

$$\mathbf{U}_p = \mathbf{U}_p^H + \mathbf{U}_p^D, \quad (2)$$

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

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

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 patho-




Fig. 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.

logical 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.




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




Fig. 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.

2.2 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). Fig. 2 shows the left putamen of a healthy volunteer and of a patient with schizophrenia, while Fig. 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 Fig. 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 intracranial 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 1 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

TABLE 1
PUTAMEN VOLUME CLASSIFICATION RATES

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

shape is really what we ought to use in the normalization.

With this in mind, our technique can be viewed as a more sophisticated version of the two features of Table 1. 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.

2.3 Organization of the Paper

The remainder of this paper is organized as follows. Section 3 motivates and 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

In this paper, we use cranium shape to normalize for non-pathological interpatient differences in the shape of brain structures. This section first motivates the approach and then describes our particular method of implementation.

3.1 Motivation

3.1.1 Brain and Cranium Shape

Different regions of the brain grow at different rates and at different times during development. The cerebellum, for example, grows later during development than does the cerebrum [23], [24]. Therefore, because cranium growth is driven by the growth of the brain underneath it [23], the cranium shape changes which are observed during neurodevelopment are at least partially caused by the heterogeneous growth of the brain.

Across a population of individuals, then, variations in brain growth will result in variations in both brain and cranium shape. Thus, there is a connection between the head shape differences observed across a population, and variations in the shapes of brain structures in that population. Because of this connection, it is possible to use cranium shape to normalize for interpatient differences in the shape of brain structures.

3.1.2 Pathology and Brain/Cranium Shape

In addition to normal variations, differences between individuals can also be caused by pathology. If this pathology occurs during neurodevelopment, it can have an influence on brain growth. Thus, as described above, both cranium growth and final cranium shape may be affected. Using

cranium shape to normalize the cranial contents for inter-patient differences can therefore result in the removal of pathological as well as nonpathological shape differences.

If the pathology begins in adulthood, however, it will affect the brain without changing the cranium. In this case, cranium shape normalization will remove only non-pathological shape differences; pathological shape differences will remain. With the interfering effects of non-pathological variation removed, the pathological variation can be more readily characterized (recall Fig. 1). The techniques developed in this paper are applicable to this second class of neurological disorder.

3.1.3 Mechanical Properties of the Brain

The above observations support the use of cranium shape in normalizing the cranial contents for interpatient differences. They do not, however, imply that the relationship between cranium shape and interior brain structure shape is linear or at all straightforward. One of the goals of the present work is to approximate this complicated and unknown relationship with a simpler, linear model and to evaluate its effectiveness.

Although their motivations have been quite different from ours, biomechanical researchers have also sought to create physical models of the human brain [25], [26], [27], [28]. The constitutive parameter values employed in these models have come from experiments on human cadavers and both living and dead animals. The complex nature of brain tissues makes it difficult to extrapolate cadaver and animal results to living humans, however. Because most of these models were created in order to study head injuries, model validation has typically been performed by computing simulations of objects hitting the head, and then comparing the simulation results to the results of analogous experiments using cadavers and animals. While some models have shown good agreement with experimentation, there is still no general consensus on either the best qualitative form of model (solid versus fluid), or on appropriate quantitative values for the brain's constitutive parameters.

Both solid and fluid models of the ICC have also been created by computer vision researchers [29], [30], [31], typically with the goal of performing some type of inter-patient warping. As in our work, the values of the mechanical properties used for these models have usually been set to those of simpler materials.

In summary, the material properties of the human brain are not fully understood [26], [28]. There are many interacting variables, and the relationship between them is still under investigation. Thus it would seem unreasonable to build models that include detailed material properties. On the other hand, it is clear that elastic deformation captures an important link between head shape, brain shape, and some disease pathology. Therefore, we would like to utilize models that include a mathematical model of the physics of deformation. Although detailed information about the underlying physical properties is still not known, even an approximate physical model of deformation could offer significant improvement in modeling and diagnosis.

3.1.4 Our Approach

In this work, we first create a physical model of the ICC. Then, using as justification the observations made in Sections 3.1.1 and 3.1.2, this model is used to normalize cranial contents for nonpathological interpatient shape variation. Finally, the remaining shape differences, which still contain the effects of pathology, are characterized statistically.

We developed this overall methodology without having available precise estimates of the brain's material properties. In the future, these parameters will become more accurately known. As these improved estimates become available, they can easily be incorporated into the general framework described in this paper.

3.2 The Finite Element Method

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. There are sound reasons for wanting to do so, as described above. 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.

The most common numerical approach for solving elastic deformation problems of this sort is the *finite element method* (FEM) [32]. 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” (^) denotes a set of FEM nodal positions or displacements, while a vector without a “hat” denotes positions or displacements at a set of nonnodal points.)

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

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

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} :

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

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

$$\mathbf{K}\hat{\mathbf{U}} = \mathbf{R} \quad (8)$$

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:

$$\mathbf{K} = \int_V \mathbf{B}^T \mathbf{C} \mathbf{B} dV, \quad (9)$$

where \mathbf{C} is the *material matrix*, which expresses the material's particular stress-strain law. See Bathe [32] for more details on setting up FEM integrals and equations.

3.3 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, (8) is diagonalized via an orthogonal transform Φ :

$$\hat{\mathbf{U}} = \Phi \tilde{\mathbf{U}}, \quad (10)$$

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

Substituting (10) into (8) and premultiplying by Φ^T yields

$$\tilde{\mathbf{K}}\tilde{\mathbf{U}} = \tilde{\mathbf{R}} \quad (11)$$

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

The optimal transformation matrix Φ is derived from the eigenvalue problem

$$\mathbf{K}\phi_i = \lambda_i \phi_i, \quad (12)$$

which, for a discretization with N nodes, has $3N$ solutions $(\lambda_1, \phi_1), (\lambda_2, \phi_2), \dots, (\lambda_{3N}, \phi_{3N})$. 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 Φ has for its columns the eigenvectors of \mathbf{K} ,

$$\Phi = [\phi_1, \phi_2, \phi_3, \dots, \phi_{3N}], \quad (13)$$

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

$$\tilde{\mathbf{K}} = \Phi^T \mathbf{K} \Phi = \Lambda = \begin{bmatrix} \lambda_1 & & & & \\ & \lambda_2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \lambda_{3N} \end{bmatrix}. \quad (14)$$

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.

3.4 FEM Model of Average ICC

3.4.1 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 [33], performing a rigid body alignment between all the extracted ICCs [34], and then averaging spatial occupancy over the aligned ICCs.

In the above procedure, the most difficult step is computing an accurate rigid body alignment between two ICCs. To accomplish this, we used the technique of Ettinger et al. [34], which we now describe. In order to align two segmented ICCs using their method, a set of 3D points is sampled from each of the two ICC surfaces. The moments of inertia of these two point sets are then used to create a coarse initial alignment. Because there is a sign ambiguity in the direction of the resultant axes, there are many possible alignments at this stage. To resolve this ambiguity, the Alignment Method [35] is employed.

This initial solution is then refined using an objective function consisting of Gaussian-smoothed interpoint distances. Finally, this refined alignment is used as the starting point of an even more detailed refinement procedure. This

second refinement procedure employs an objective function that uses a rectified least squares measure. Ettinger et al. [34] argue that their first refinement procedure is good at generating an alignment that is close to the global minimum, while their second refinement procedure is good at precisely locating a minimum that is initially nearby.

Nonrigid methods such as trainable snakes [56] were also considered for constructing the average ICC. Unfortunately, current implementations of such methods rely on consistent sampling and labeling of point features across the entire data set. Furthermore, the point labeling imposes a parameterization on the data. If different feature points are present in different views, or if there are very different sampling densities, then the resulting models will differ even if the object's pose and shape are identical.

3.4.2 FEM Model

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:

$$\mathbf{X}_A^{ICC} = [x_1, y_1, z_1, \dots, x_M, y_M, z_M]^T. \quad (15)$$

An initial coarse alignment between this surface data and the deformable model is then found by computing the data's moments of inertia. 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:

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

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 (8) is therefore constructed as

$$\mathbf{R}_A = \mathbf{H}^T \mathbf{F}_A, \quad (17)$$

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:

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

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

$$\hat{\mathbf{X}}_A^{ICC} = \hat{\mathbf{X}} + \hat{\mathbf{U}}_A^{ICC}, \quad (19)$$

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

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

$$\tilde{\mathbf{U}}_A^{ICC} = \Phi^T \hat{\mathbf{U}}_A^{ICC}, \quad (20)$$

where the modal coefficient vector $\tilde{\mathbf{U}}_A^{ICC}$ specifies how much of each deformation mode is contained in the shape




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

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 (11). As mentioned earlier, because the system of equations is decoupled in the modal coordinate system, using (11) directly is much faster.

Fig. 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 [36], [60] 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.

3.5 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:

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

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

$$\hat{\mathbf{X}}_p^{ICC} = \hat{\mathbf{X}} + \hat{\mathbf{U}}_p^{ICC}. \quad (22)$$

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

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.

3.6 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 [37]. Because this will produce noninteger 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.

3.6.1 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 (4), using the nodal positions of the average ICC:

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

Note, however, that (23) must be inverted—given \mathbf{x} , we need to find \mathbf{r} .

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

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

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

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

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

$$\mathbf{J}^k(\mathbf{x}^{k+1} - \mathbf{x}^k) + \mathbf{g}^k = 0, \quad (26)$$

where \mathbf{J} is the Jacobian matrix defined in (7), and k represents iteration [38], [39].

3.6.2 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 (4), but this time we use it directly, with the known \mathbf{r} and the ICC nodal positions of patient p :

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

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.

3.6.3 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 (4), the interpolation and warping is done by combining (5) and (10):

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

To further increase computational efficiency, the polynomial deformations of (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:

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

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

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

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

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

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

Equations (30) and (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.




Fig. 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).

3.7 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 Fig. 5. Figs. 5a and 5b show the average ICC and ventricles, while Figs. 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 Fig. 5e. As can be seen, these warped ventricles are more similar to the average ventricles in Fig. 5b.

4 CHARACTERIZATION OF DISEASE STATES

4.1 Shape Representation

As described in (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, unwrapped 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 dis-

placement 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, 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 (2).

We have therefore met our original goal of separating the total deformation into its pathological and non-pathological 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 2D outline or 3D 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

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

where $\bar{\mathbf{U}}^D$ 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 Ψ . Any patient's displacement vector can be written as a linear combination of the principal components:

$$\mathbf{U}_p^D = \Psi \mathbf{b}_p, \quad (33)$$

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.




Fig. 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).

Designing a Gaussian quadratic classifier requires estimating the 6×6 covariance matrix Σ_i of each of the C classes present. However, with only between seven and 13 members per class, the Σ_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 $\Sigma_1 = \Sigma_2 = \dots = \Sigma_C = \Sigma$. Thus, only the overall Σ 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 classes¹:

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

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 (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 [40].

In order to separate the training stage from the classification stage, we use the *leave-one-out* [41] 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.

In summary, the steps of our classification for each of the 25 patients are:

- 1) Use the remaining 24 patients to calculate Σ and the C class means \mathbf{m}_i .

- 2) Compute the C discriminant functions g_i (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.² 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.

Fig. 6 shows the first two putamen principal components. The first mode (Figs. 6a-6e) is a contrast between the size of the top of the putamen and the size of the bottom, while the second (Figs. 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 2 shows the results, along with the classification rate using just putamen volume. A 12 percent improvement in the classification rate occurs when using the putamen principal components instead of just volume.

To test the robustness of this classification result, we varied the number of principal components input into the

1. Equation (34) assumes that the prior probabilities of all classes are equal.

2. Details of the MR acquisition, along with the clinical diagnostic procedures, can be found in [7] and [17].




Fig. 7. Classification rates obtained when varying the number of putamen principal components input into a Gaussian linear classifier.

Gaussian linear classifier. Fig. 7 shows the result. As can be seen, there is a region of stability that begins when the top $\frac{1}{4}$ (6) principal components are used. The classification rate then degrades as the higher order principal components, which contain a high degree of noise, are included in the analysis.

5.2 Ventricular Disorders

Nine patients with Alzheimer's disease, seven patients with normal-pressure hydrocephalus, and nine healthy control subjects, all matched for age, underwent an MR brain scan.³ As part of a previous volumetric study [22], the ventricles were segmented using a semiautomatic procedure. Using precisely the same steps that were applied to the putamen, we estimated the principal components of the ventricular data set.

3. Details of the MR acquisition, along with the clinical diagnostic procedures, can be found in [22].

TABLE 2
PUTAMEN CLASSIFICATION RATES

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

TABLE 3
VENTRICLE CLASSIFICATION RATES

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

Fig. 8 shows two of the principal components. The first mode (Figs. 8a-e) is just a measure of the overall size of the ventricles, while the second (Figs. 8f-j) is the degree of extension of the posterior horns of the lateral ventricles.

Table 3 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, there is an improvement seen when using principal components instead of just volume.

Fig. 9 shows the result obtained when the number of principal components input to the classifier is varied. As can be seen, the region starting at the top $\frac{1}{4}$ (6) principal components is relatively stable when compared to the other regions. This region is, however, less stable than it was for the putamen database (Fig. 7). This and other difficulties with the ventricle database are discussed in the next section.

6 DISCUSSION

6.1 Head Shape Normalization

To test the effects of head shape normalization, we computed the putamen principal components *without* first warping all patients' cranial contents to the same model.




Fig. 8. (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).




Fig. 9. Classification rates obtained when varying the number of ventricle principal components input into a Gaussian linear classifier.

For the putamen data set, this dropped the classification rate from 72 percent to 64 percent, 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. Fig. 10a shows the projections onto the top

two modes of the unwarped data, while Fig. 10b shows the projections onto the top two modes computed after warping. Note that b_1 and b_2 of Fig. 10b are the first two components of the projection vector \mathbf{b} of (33), and that the eigenvectors onto which they are projected were shown earlier in Fig. 6. Similarly, a_1 and a_2 of Fig. 10a 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, Fig. 10a shows less class separability than Fig. 10b. 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 Fig. 10c. 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 percent. Fig. 11 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 Figs. 11a and 11b.

This result, along with the lack of robustness with respect to the number of principal components input into the classifier observed in Section 5.2, has several possible causes. One possible explanation for the lack of improvement when normalizing for head shape is the already high classification rate (88 percent) obtainable without removing

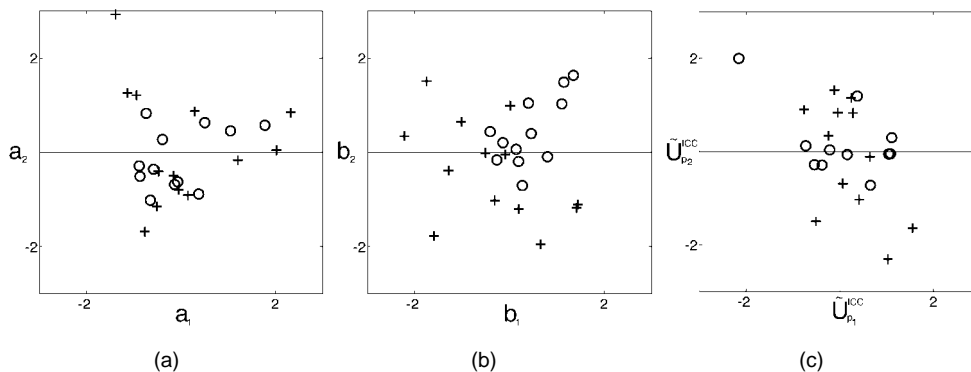


Fig. 10. Projections onto modes computed from the schizophrenia database. Each schizophrenic patient is denoted by a +, and each healthy volunteer is denoted by a \circ . (a) Putamen principal components, computed from the original data. (b) Putamen principal components, computed after first normalizing for head shape. (c) ICC physical deformation modes.

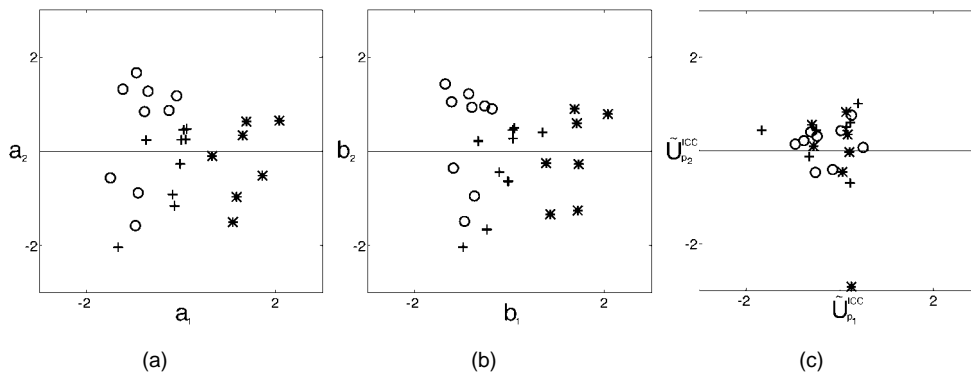


Fig. 11. Projections onto modes computed from the ventricular database. Each Alzheimer’s patient is denoted by a +, each normal-pressure hydrocephalus patient by an *, and each healthy volunteer by a \circ . (a) Ventricle principal components, computed from the original data. (b) Ventricle principal components, computed after first normalizing for head shape. (c) ICC physical deformation modes.

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's, but interior structures do not necessarily scale by precisely the same amount [42]. 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.

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 (8) and its solution

$$\hat{\mathbf{U}} = \mathbf{K}^{-1}\mathbf{R} \quad (35)$$

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

$$\Sigma_{\hat{\mathbf{U}}} = \mathbf{K}^{-1}\Sigma_{\mathbf{R}}(\mathbf{K}^{-1})^T \quad (36)$$

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

$$\Sigma_{\hat{\mathbf{U}}} = \mathbf{K}^{-1}\Sigma_{\mathbf{R}}\mathbf{K}^{-1}. \quad (37)$$

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

$$\Sigma_{\hat{\mathbf{U}}} = \sigma^2\mathbf{K}^{-2}. \quad (38)$$

We can form the same covariance matrix $\mathbf{S} \approx \Sigma_{\hat{\mathbf{U}}}$ from a set of observations of $\hat{\mathbf{U}}$, then use (38) to obtain the estimate

$$\mathbf{S} \approx \sigma^2\mathbf{K}^{-2}. \quad (39)$$

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 3.3 that diagonalized \mathbf{K} , we have

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

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 Φ also diagonalizes \mathbf{S} , which implies that Φ 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. For the data sets used in this paper, it is interesting to note that some of the physical ICC modes had similarities to the statistically derived modes of the interior brain structures.

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 relate each patient's ICC surface to a model of an average ICC. Equations (35) to (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 [29] 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. [43], [44] 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. [45] 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 [46], [47]. 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 [48] and their variants. Cohen [49] augmented the original snake formulation with a balloon force to help it avoid local minima. Staib and Duncan [50] 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's rule in order to find the best set of parameter values. Székely et al. [51] extended this approach to 3D using the spherical harmonic technique developed by Brechbühler [52], 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 [50] and Székely et al. [51] 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 [53] required corresponding landmarks, more recent efforts [54] 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 [55] have used principal components to describe facial variation and have been able to use this approach to reliably recognize people's faces. Cootes et al. [56] used principal components to experimentally describe the modes of variation inherent in a training set of 2D heart images. Hill et al. [57] extended this technique to 3D and analyzed the cerebral ventricles for purposes of segmentation and representation.

Along with ourselves [58], 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 [59] have also considered both physical and statistical shape models, in the context of representing and recognizing objects from their 2D 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 2D 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 intracranial 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 principal 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. This work was supported in part by NIH grant P01 CA67165-02.

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*, vol. 116, pp. 111-115, July 1975.
- [2] C.P. Hughes and M. Gado, "Computed Tomography and Aging of the Brain," *Radiology*, vol. 139, pp. 391-396, May 1981.
- [3] T. El Gammal, M.B. Allen, Jr., B.S. Brooks, and E.K. Mark, "MR Evaluation of Hydrocephalus," *Amer. J. Neuroradiology*, vol. 8, pp. 591-597, July/Aug. 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*, vol. 31, pp. 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," *Amer. J. Neuroradiology*, vol. 10, pp. 371-376, Mar./Apr. 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," *Amer. J. Neuroradiology*, vol. 9, pp. 1,181-1,187, Nov./Dec. 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," *New England J. of Medicine*, vol. 327, no. 9, pp. 604-612, Aug. 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*, vol. 43, pp. 719-725, Apr. 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*, vol. 43, pp. 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*, vol. 43, pp. 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*, vol. 43, pp. 1,083-1,087, June 1993.
- [12] E.H. Aylward, J.D. Henderer, J.C. McArthur, P.D. Bretschneider, G.J. Harris, P.E. Barta, and G.D. Pearlson, "Reduced Basal Ganglia Volumes in HIV-1-Associated Dementia: Results From Quantitative Neuroimaging," *Neurology*, vol. 43, pp. 2,099-2,104, Oct. 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*, vol. 43, pp. 2,117-2,124, Oct. 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*, vol. 43, pp. 2,380-2,382, Nov. 1993.
- [15] A.M. Murro, Y.D. Park, D.W. King, B.B. Gallagher, J.R. Smith, F. Yaghami, 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*, vol. 43, pp. 2,531-2,533, Dec. 1993.
- [16] H.S. Soinenen, 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*, vol. 44, pp. 1,660-1,668, Sept. 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. McCarterley, "Caudate, Putamen, and Globus Pallidus Volume in Schizophrenia: A Quantitative MRI Study," unpublished manuscript, 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," *J. Computer Assisted Tomography*, vol. 11, no. 6, pp. 923-931, Nov./Dec. 1987.
- [19] A.P. Pentland and S. Sclaroff, "Closed-Form Solutions for Physically Based Shape Modeling and Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 715-729, July 1991.
- [20] S. Sclaroff and A.P. Pentland, "On Modal Modeling for Medical Images: Underconstrained Shape Description and Data Compression," *CVPR Workshop Biomedical Image Analysis*, June 1994.
- [21] T.F. Cootes, "Combining Point Distribution Models With Shape Models Based on Finite Element Analysis," *Proc. British Machine Vision Conf.*, 1994.
- [22] M. Matsumae, R. Kikinis, I. Mórocz, A.V. Lorenzo, M.S. Albert, P.M. Black, and F.A. Jolesz, "Intracranial Compartment Volumes in Patients With Enlarged Ventricles Assessed by Magnetic Resonance-Based Image Processing," *J. Neurosurgery*, vol. 84, pp. 972-981, June 1996.
- [23] J.M. Tanner, *Fetus Into Man: Physical Growth From Conception to Maturity*. Cambridge, Mass.: Harvard Univ. Press, 1990.
- [24] J.J. Volpe, *Neurology of the Newborn*. W.B. Saunders Company, 1995, 3rd ed.
- [25] M. Chan, C. Ward, D. Schneider, and S. Adams, "Relative Importance of Skull Deformation," *Proc. Biomechanics Symp.*, pp. 157-160, 1981.
- [26] C. Chu, M. Lin, and M. Lee, "Finite Element Analysis of Cerebral Contusion," *J. Biomechanics*, vol. 27, no. 2, pp. 187-194, 1994.
- [27] W. Goldsmith, "Biomechanics, Its Foundations and Objectives," Y.C. Fung, N. Perrone, and M. Anliker, eds., *Biomechanics of Head Injury*. Englewood Cliffs, N.J.: Prentice Hall, Inc., 1972.
- [28] J.S. Ruan, T. Khalil, and A.I. King, "Dynamic Response of the Human Head to Impact by Three-Dimensional Finite Element Analysis," *J. Biomechanical Eng.*, vol. 116, pp. 44-50, 1994.
- [29] R. Bajcsy and S. Kovacic, "Multiresolution Elastic Matching," *Computer Vision Graphics Image Processing*, vol. 46, pp. 1-21, 1989.
- [30] C.E. Christensen, M.I. Miller, and M. Vannier, "A 3D Deformable Magnetic Resonance Textbook Based on Elasticity," *Proc. AAAI Symp. Applications of Computer Vision in Medical Image Processing*, Mar. 1994.
- [31] D.L. Collins, T.M. Peters, W. Dai, and A.C. Evans, "Model Based Segmentation of Individual Brain Structures From MRI Data," *SPIE: Visualization in Biomedical Computing*, 1992, pp. 10-23.
- [32] K. Bathe, *Finite Element Procedures in Engineering Analysis*. Englewood Cliffs, N.J.: Prentice Hall, Inc., 1982.
- [33] H.E. Cline, W.E. Lorensen, R. Kikinis, and F.A. Jolesz, "3D Segmentation of MR Images of the Head Using Probability and Connectivity," *J. Computer Assisted Tomography*, vol. 14, no. 6, pp. 1,037-1,045, 1990.
- [34] G. Ettinger, E. Grimson, and T. Lozano-Perez, "Automatic Registration for Multiple Sclerosis Change Detection," *CVPR Workshop Biomedical Image Analysis*, June 1994.
- [35] D. Huttenlocher and S. Ullman, "Recognizing Solid Objects by Alignment With an Image," *Int'l J. Computer Vision*, vol. 5, no. 2, pp. 195-212, 1992.
- [36] S. Sclaroff and A.P. Pentland, "Modal Matching for Correspondence and Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 6, pp. 545-561, June 1995.
- [37] G. Wolberg, *Digital Image Warping*. Los Alamitos, Calif.: IEEE CS Press, 1990.
- [38] G. Strang, *Introduction to Applied Mathematics*. Wellesley-Cambridge Press, 1986.
- [39] I.A. Essa, S. Sclaroff, and A.P. Pentland, "Physically-Based Modeling for Graphics and Vision," R. Martin, ed., *Directions in Geometric Computing 1992*. Information Geometers, 1992.
- [40] C.W. Therrien, *Decision, Estimation, and Classification: An Introduction to Pattern Recognition and Related Topics*. New York: John Wiley & Sons, 1989.
- [41] K. Fukunaga, *Introduction to Statistical Pattern Recognition*. New York: Academic Press, 1990.
- [42] M. Matsumae, R. Kikinis, I. Mórocz, A.V. Lorenzo, T. Sándor, M.S. Albert, P.M. Black, and F.A. Jolesz, "Age-Related Changes in Intracranial Compartment Volumes in Normal Adults Assessed by Magnetic Resonance Imaging," *J. Neurosurgery*, vol. 84, pp. 982-991, June 1996.
- [43] D.L. Collins, P. Neelin, T.M. Peters, and A.C. Evans, "Automatic 3D Intersubject Registration of MR Volumetric Data in Standardized Talairach Space," *J. Computer Assisted Tomography*, vol. 18, no. 2, pp. 192-205, 1994.
- [44] 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," *SPIE: Visualization in Biomedical Computing*, 1994, pp. 180-190.
- [45] C.E. Christensen, R.D. Rabbitt, and M.I. Miller, "3D Brain Mapping Using a Deformable Neuroanatomy," *Physics in Medicine and Biology*, vol. 39, pp. 609-618, 1994.
- [46] R. Szeliski, *Bayesian Modeling of Uncertainty in Low-Level Vision*. Kluwer Academic Publishers, 1989.
- [47] T.E. Boult, S.D. Fenster, and T. O'Donnell, "Physics in a Fantasy World vs Robust Statistical Estimation," *Proc. NSF Workshop 3D Object Recognition*, New York, Nov. 1994.
- [48] D. Terzopoulos, A. Witkin, and M. Kass, "Symmetry-Seeking Models and 3D Object Reconstruction," *Int'l J. Computer Vision*, vol. 1, pp. 211-221, 1987.
- [49] L. Cohen, "On Active Contour Models and Balloons," *Computer Vision, Graphics and Image Processing: Image Understanding*, vol. 53, no. 2, pp. 211-218, 1991.
- [50] L.H. Staib and J.S. Duncan, "Boundary Finding With Parametrically Deformable Models," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 11, pp. 1,061-1,075, Nov. 1992.
- [51] 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," unpublished manuscript, 1995.
- [52] 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.
- [53] F.L. Bookstein, "Principal Warps: Thin-Plate Splines and the Decomposition of Deformations," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 6, pp. 567-585, June 1989.
- [54] F.L. Bookstein and W.D.K. Green, "The Biometrics of Landmarks and Edgels: A New Geometry of Prior Knowledge for Medical Image Understanding," *Proc. AAAI Symp. Applications of Computer Vision in Medical Image Processing*, pp. 134-137, Mar. 1994.
- [55] M. Turk and A.P. Pentland, "Eigenfaces for Recognition," *J. Cognitive Neuroscience*, vol. 3, no. 1, pp. 71-86, 1991.

- [56] T.F. Cootes, D.H. Cooper, C.J. Taylor, and J. Graham, "Trainable Method of Parametric Shape Description," *Image and Vision Computing*, vol. 10, no. 5, pp. 289-294, June 1992.
- [57] A. Hill, T.F. Cootes, and C.J. Taylor, "A Generic System for Image Interpretation Using Flexible Templates," *Proc. British Machine Vision Conf.*, pp. 276-285, 1992.
- [58] J. Martin, A.P. Pentland, and R. Kikinis, "Shape Analysis of Brain Structures Using Physical and Experimental Modes," *Proc. CVPR*, pp. 752-755, 1994.
- [59] S.C. Zhu and A.L. Yuille, "Forms: A Flexible Object Recognition and Modelling System," Harvard Robotics Laboratory, no. 94-1, 1993.
- [60] Sclaroff, S., "Deformable Prototypes for Encoding Shape Categories in Image Databases," *Pattern Recognition*, vol. 30, no. 4, pp. 627-642, Apr. 1997.



John Martin received the BS degree in mathematics and the MS degree in applied mathematics from Rensselaer Polytechnic Institute in 1984. From 1984 to 1988, he worked as a software engineer for Hughes Aircraft Company. He earned the PhD degree from the Massachusetts Institute of Technology in 1995, studying radiological sciences and doing research in medical computer vision. From 1996 to 1997, Dr. Martin worked at Visualization Technology, Inc., where he helped develop image-guided surgery systems.

Since June 1997, he has been employed at Millennium Pharmaceuticals, Inc., where he applies signal processing and machine learning techniques to genomics. His current research interests include pattern recognition, genetics, and the modeling of biological systems.



Alex Pentland received his PhD from the Massachusetts Institute of Technology in 1982. He is currently the academic head of the MIT Media Laboratory. He is also the Toshiba Professor of Media Arts and Sciences, an endowed chair last held by Marvin Minsky. He has worked at SRI's AI Center and as a lecturer at Stanford University, winning the Distinguished Lecturer award in 1986. In 1987, he returned to MIT to found the Perceptual Computing Section of the Media Laboratory, a group that now includes more than

50 researchers in computer vision, graphics, speech, music, and human-machine interaction. He has done research in wearable computing, human-machine interface, computer graphics, artificial intelligence, machine and human vision and has published more than 180 scientific articles in these areas. His most recent research focus is understanding human behavior in video, including face, expression, gesture, and intention recognition, as described in the April 1996 issue of *Scientific American*. He has won awards from the AAAI for his research into fractals, from the IEEE for his research into face recognition, and from *Ars Electronica* for his work in computer vision interfaces with virtual environments. *Newsweek* magazine has recently named him one of the 100 Americans most likely to shape the next century.



Stan Sclaroff received the MS degree and PhD degree from the Massachusetts Institute of Technology in 1991 and 1995, respectively. He is currently an assistant professor in the Computer Science Department at Boston University, where he has founded the Image and Video Computing Group. In 1995, he received a Young Investigator Award from the Office of Naval Research and a Faculty Early Career Development Award from the National Science Foundation.

From 1989 to 1994, he was a research assistant in the Vision and Modeling Group at the MIT Media Laboratory. Prior to that, he worked as a senior software engineer in the solids modeling and computer graphics groups at Schlumberger Technologies, CAD/CAM Division.



Ron Kikinis received his MD from the University of Zurich, Switzerland, in 1982. He has worked as a researcher at the ETH in Zurich and as a resident at the University Hospital in Zurich. He is currently the director of the Surgical Planning Laboratory of the Department of Radiology, Brigham and Women's Hospital and Harvard Medical School, an associate professor of radiology at Harvard Medical School, as well as an adjunct professor of biomedical engineering at Boston University. His research interests include

the development of clinical applications for image processing, computer vision, and interactive rendering methods. He is currently concentrating on developing fully automated segmentation methods and introducing computer graphics into the operating room. He is the author of 52 peer-reviewed articles.