

Shortened version to appear in IEEE CVPR Seattle WA June 1994

Shape Analysis of Brain Structures Using Physical and Experimental Modes

John Martin, Alex Pentland, 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, kikinis@bwh.harvard.edu

Abstract

This paper presents 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. By using the equivalence of physically-based and probability-based formulations of deformation, the total shape deformation is decomposed into physical modes of variation obtained from finite element analysis, and experimental modes of variation obtained from sample data using principal component analysis.

This mathematical model is used to classify diseases that affect the shape of the ventricular system of the brain. Because ventricular shape 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 ventricular shape variation. This allows us to experimentally characterize modes of variation that are indicative of disease processes.

Our technique is applied to magnetic resonance images of the brains of individuals with 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. Hydrocephalus, for example, affects the morphology of the ventricular system. These changes have been studied for several decades, using both postmortem and 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.

Because of both genetic and environmental factors, however, biological structures have a large normal range of variation. The onset of a disease introduces further changes in morphology, with each particular disease causing its own type of shape changes and range of variation in these changes. Therefore to properly study the pathological deformations, we must first account for the large normal variation in biological shape.

For structures inside the cranium, part of this normal shape variability is due to the different head shapes seen across individuals. These macroscopic shape effects interfere with the analysis of local shape deformations caused by disease. For example, the slight skewness of a patient's brain ventricles may be

due solely to a slight skewness in his head shape. However if a disease was thought to cause ventricular skewness, this patient might be misdiagnosed if his overall head shape was not taken into account.

What is needed then is a method that separates out these two types of deformations, allowing just the disease deformations to be analyzed. In this paper we create a mathematical framework that (1) separates out disease deformation from head shape deformation, and (2) allows us to represent the deformations caused by disease in an intuitive manner. This is accomplished 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. Because detailed physical models of neurological diseases do not exist, we turn to statistical techniques to examine this data.

By using the technique of modal analysis when doing the physical modeling, and the technique of principal component analysis when doing the statistics, it is possible to express both the deformations caused by head shape and the deformations caused by diseases in terms of modes of variation. These modes represent unique, natural coordinates in which to express the shape and deformation of brain structures. The modes can be (1) displayed to show how the structures deform due to both head shape and pathology, and (2) used in pattern recognition algorithms to classify diseases based on shape changes.

We have selected the ventricles of the brain to study disease and head shape deformations. Figure 1 shows the ventricles of a healthy volunteer, a patient with Alzheimer’s disease (AD), and a patient with normal-pressure hydrocephalus (NPH). Both of the disorders cause the ventricles to enlarge, but in different amounts and in different ways.

As a reference point for the methodology to be developed in the following sections, we consider here possible ways of classifying patients into the three classes shown in Figure 1. Given a data set consisting of samples from each of these three categories, the most straightforward classification procedure would be to use just one feature, ventricular volume. However since a person with a larger head will tend to have larger ventricles, even if he is healthy he may be misclassified into one of the disease groups. Therefore the second procedure to try would be to normalize each person’s ventricular volume by his head volume.

Using each of the above two features, we ran a Gaussian quadratic classifier on a data set consisting

Feature	# Correct
Volume	22
Normalized volume	20

Table 1: Number of patients correctly classified (out of 25), using ventricular volume and ventricular volume normalized by overall ICC volume.

of 25 people, approximately evenly divided into normals, AD patients, and NPH patients. Table 1 shows the results.¹

Contrary to these results, one would expect that the classification should improve when normalizing for overall ICC volume. While the decrease in classification 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.

With this in mind, we have developed a more sophisticated version of the two features of Table 1. Instead of using just ventricular volume, a principal component analysis of the ventricular 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. Later sections of the paper describe how this is accomplished using experimental and physical deformation modes.

Although from a medical viewpoint having access to both AD and NPH patients makes this a fairly unique data set, 25 patients (divided into three groups) is not enough to claim statistical significance in the experiments we perform. Instead the main contributions of this work are (1) to suggest a procedure that intuitively make sense, (2) to develop a consistent mathematical formulation for this procedure, and (3) to show that the data we do have indicates that the procedure is working. In the future we hope to apply our method to larger data sets.

This paper is organized as follows. Section 2 reviews the relevant literature, comparing our work to that of several other researchers. In Section 3 we present the mathematical model describing the physical and experimental modes of deformation. Section 4

¹It should be noted that because of our small sample size, we decided not to hold out the patient being classified when computing the classification boundaries. This has the effect of inflating the classification rate. Therefore it is the difference between the classification numbers reported throughout this paper that is important, not their absolute values.

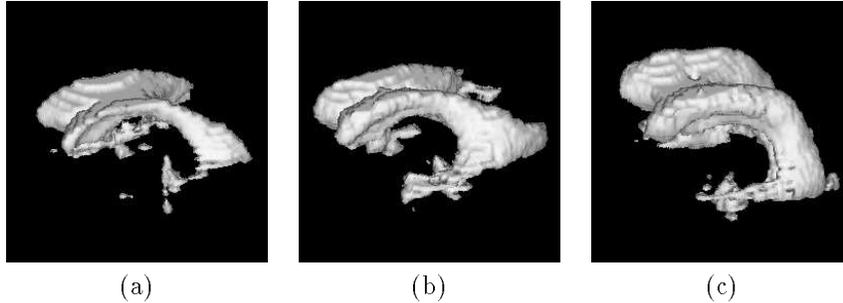


Figure 1: Reconstructions of the ventricular system of the human brain created from MR images. (a) Normal, healthy adult. The ventricles are not very pronounced. (b) Alzheimer’s disease causes an enlargement of the ventricles. (c) Normal-pressure hydrocephalus causes an even greater enlargement of the ventricles, along with an overall puffy appearance.

then applies the mathematical model to brain structures. The ventricular modes of variation are examined qualitatively in Section 5, while Section 6 presents our more quantitative experimental results. Section 7 discusses the unresolved issues pertaining to this work. Section 8 summarizes the work done.

2 Related work

Terzopoulos *et al.* [1] used a physical model with a very large number of degrees of freedom to perform dynamic shape fitting. Pentland and Sclaroff [2] employed a much smaller number of degrees of freedom, representing shape in terms of an object’s physical deformation modes. Instead of using the modes of a particular object, Bookstein [3] described shape deformation in terms of the physical deformation modes of an infinite thin plate, applying it primarily to the biological domain. Collins *et al.* [4] used this thin plate model to calculate brain deformation. This deformation was then applied to internal structures of a predefined anatomical atlas in order to improve segmentation.

Instead of physically modeling the structure under study, researchers have also sought to approximate the modes of variation through experimental observation. Turk and Pentland [5] 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.* [6] used principal components to experimentally describe the modes of variation inherent in a training set of 2D heart images. Hill *et al.* [7] extended this technique to 3D and analyzed the ventricles of the brain for purposes of segmentation and representation.

Our method differs from the above work in the following respects. First, all of the above techniques

use *either* physical modeling or experimental observation in order to represent shape deformation. In our method, we combine *both* physical and experimental techniques in order to describe shape variation. Second, in addition to representing the modes of variation in healthy brain structures, we also investigate the modes of variation caused by different disease processes. Finally, we use these modes to classify patients into different disease categories.

3 Mathematical model

3.1 Introduction

The basic conceptual framework in which we are placing this problem is as follows. The differences seen across people in the shape of any particular brain structure are a combination of three separate effects: overall head shape, any other normal differences not accounted for by head shape, and any pathological differences due to disease processes. We cannot do anything about the normal differences not accounted for by head shape. The pathological differences we want to leave in for the classification. The differences due to head shape we can take out.

Before discussing the application area in greater detail, we first develop a consistent mathematical framework. This is done in detail below. A brief summary of what follows is given now. First, we note that by creating an analytical model of an object we can obtain, by using the technique of *modal analysis*, a shape description in terms of the free vibration modes of the object. On the other hand, by gathering sample data of the object instead of characterizing it analytically, we can use the statistical technique of *principal component analysis* to find the directions of greatest variability in the data. This decomposition also allows for

a shape description, but this time it is in terms of the principal components of the data sets. We will show how an object’s shape, for which we have both a physical model and some experimental data, can be written in terms of a combination of the physical modes obtained through modeling and the statistical modes obtained through experimentation.

3.2 Physical modes

3.2.1 Physical modeling

The derivation of the physical modes begins with the creation of a physical model of the object under study. In our case, we model the object as a linear elastic material, and then set up equations describing its dynamics. The solution to these equations gives the elastic deformation that the object undergoes due to applied forces.

The most common numerical approach for solving problems of this sort is the *finite element method* (FEM). 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 uniformly over the body as a function of its constitutive parameters.

In general, the polynomial shape function for each element is written in vector form as:

$$\mathbf{u}(\mathbf{x}) = \mathbf{H}(\mathbf{x})\mathbf{U} \quad (1)$$

where \mathbf{H} is the interpolation matrix, \mathbf{x} is the local coordinate of a point in the element where we want to know the displacement, and \mathbf{U} denotes a vector of displacement components at each element node.

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:

$$\epsilon(\mathbf{x}) = \mathbf{B}(\mathbf{x})\mathbf{U} \quad (2)$$

where \mathbf{B} is the strain displacement matrix. The rows of \mathbf{B} are obtained by appropriately differentiating and combining rows of the element interpolation matrix \mathbf{H} .

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

$$\mathbf{M}\ddot{\mathbf{U}} + \mathbf{K}\mathbf{U} = \mathbf{R}, \quad (3)$$

where \mathbf{R} is the load vector whose entries are external forces acting on the nodes, and \mathbf{M} and \mathbf{K} are the element mass and stiffness matrices, respectively. \mathbf{M} and \mathbf{K} are computed directly from the interpolation and strain displacement matrices by integrating over the object’s volume:

$$\mathbf{M} = \int_V \rho \mathbf{H}^T \mathbf{H} dV \quad \text{and} \quad \mathbf{K} = \int_V \mathbf{B}^T \mathbf{C} \mathbf{B} dV, \quad (4)$$

where ρ is the mass density, and \mathbf{C} is the *material matrix* which expresses the material’s particular stress-strain law. See Bathe [8] for more details on setting up FEM integrals and equations.

If only the final equilibrium state is required, then we can just solve the static equilibrium equation

$$\mathbf{K}\mathbf{U} = \mathbf{R}. \quad (5)$$

This can also be formulated in terms of the energy of the system:

$$E(\mathbf{U}) = \frac{1}{2} \mathbf{U}^T \mathbf{K} \mathbf{U} + \mathbf{P}(\mathbf{U}) \quad (6)$$

where $\mathbf{P}(\mathbf{U})$ is the potential field that gives rise to the nodal force vector \mathbf{R} . The minimization of $E(\mathbf{U})$ results in Equation 5.

3.2.2 Modal analysis

The technique of modal analysis consists of diagonalizing Equation 5 by applying an orthogonal transform Φ :

$$\mathbf{U} = \Phi \tilde{\mathbf{U}} \quad (7)$$

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 Equation 7 into Equation 5 and pre-multiplying by Φ^T yields

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

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

$$\mathbf{K} \phi_i = \lambda_i \phi_i \quad (9)$$

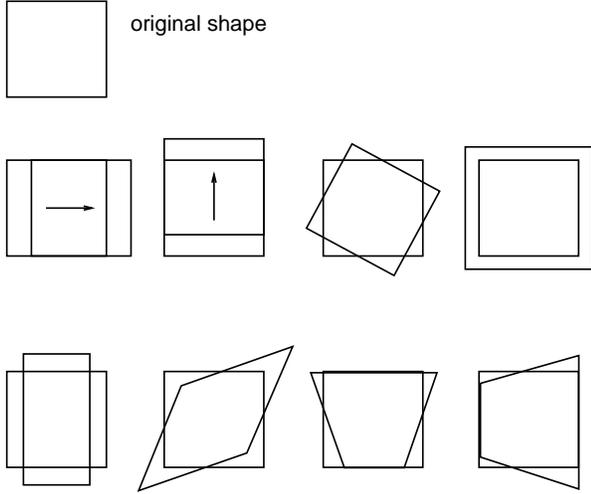


Figure 2: Physical modes of a square.

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. In our work we will assume that $\mathbf{M} = m\mathbf{I}$ (where m is the amount of mass assigned to each node), as the voxels are evenly spaced, the material is of roughly uniform density, and this assumption simplifies our computation.

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

$$\Phi = [\phi_1, \phi_2, \phi_3, \dots, \phi_{3N}], \quad (10)$$

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

$$\tilde{\mathbf{K}} = \Phi^T \mathbf{K} \Phi = \mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & & & & \\ & \lambda_2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \lambda_{3N} \end{bmatrix} \quad (11)$$

Using this diagonalization we are able to write any arbitrary displacement of the nodes as a linear combination of the physical modes:

$$\mathbf{U} = \Phi \mathbf{a} = \sum_{i=1}^{3N} \Phi_i a_i \quad (12)$$

where \mathbf{a} is a vector of the modal coefficients a_i .

This is illustrated by the two-dimensional example shown in Figure 2. By placing one node at each corner of the square and assigning material properties,

a stiffness matrix \mathbf{K} can be constructed. Next, the modal transformation is applied to the system. The resulting eight free vibration modes (the eigenvectors of \mathbf{K}) are shown in the figure.

Because the stiffness matrix has been diagonalized, the resulting system of equations is decoupled and therefore computationally much easier to handle. 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.

Software that uses the FEM and modal analysis to recover and describe shapes is available from whitechapel.media.mit.edu in the file `/u/ftp/misc/modal.tar.Z`.

3.3 Experimental modes

Instead of modeling the system as an elastic material, we can instead assume nothing about it and just collect data samples of the displacements of each node. Here each node is treated as a random variable, and after collecting P samples of deformed objects, we can form the sample covariance matrix of these measurements:

$$\mathbf{S} = \frac{1}{P-1} \sum_{p=1}^P (\mathbf{U} - \bar{\mathbf{U}})(\mathbf{U} - \bar{\mathbf{U}})^T. \quad (13)$$

Diagonalizing \mathbf{S} by performing a principal component analysis gives the experimental modes of variation, which we assemble as the columns in a matrix Ψ . Any arbitrary displacement of the nodes can now be written as a linear combination of these experimental modes²:

$$\mathbf{U} = \Psi \mathbf{b} = \sum_{i=1}^{3N} \Psi_i b_i. \quad (14)$$

In this case it is the modes whose associated eigenvalues are *large* that we are typically most interested in. This is because the directions of highest variability in the data are the eigenvectors of the largest eigenvalues, and in fact these are the most reliably estimated eigenvectors.

²This assumes that the training set is large enough to represent all the possible variation in the new samples. Variation orthogonal to the space spanned by the training set will not be represented.

Software that calculates and analyzes experimental modes is available from whitechapel.media.mit.edu in the file `/u/ftp/misc/facerecognition.tar.Z`.

3.4 The connection

The connection between physical and experimental modes comes from the close relationship between mechanical and probabilistic prior models [9][10]. The mechanical viewpoint is the one we have used above, modeling the elastic field by a stiffness matrix \mathbf{K} and then minimizing the deformation energy which is a function both of the displacements and of \mathbf{K} (Equation 6). As mentioned, given \mathbf{K} we can analytically compute the modes of variation, independent of any assumptions about $\mathbf{P}(\mathbf{U})$ and \mathbf{R} .

We can also interpret Equation 6 and its solution

$$\mathbf{U} = \mathbf{K}^{-1}\mathbf{R} \quad (15)$$

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

$$\boldsymbol{\Sigma}_U = \mathbf{K}^{-1}\boldsymbol{\Sigma}_R(\mathbf{K}^{-1})^T \quad (16)$$

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

$$\boldsymbol{\Sigma}_U = \mathbf{K}^{-1}\boldsymbol{\Sigma}_R\mathbf{K}^{-1}. \quad (17)$$

If the elements of \mathbf{R} are independent and have unit variance, then

$$\boldsymbol{\Sigma}_U = \mathbf{K}^{-2}. \quad (18)$$

Therefore from a set of observations of \mathbf{U} , we can form the sample covariance matrix \mathbf{S} as an estimate of the true covariance $\boldsymbol{\Sigma}_U$, then use Equation 18 to write

$$\mathbf{S} \approx \mathbf{K}^{-2}. \quad (19)$$

By collecting samples of \mathbf{U} and using Equation 19, then, we can approximate the stiffness matrix \mathbf{K} .

This connection leads to several useful observations. First, by using a physical model, we are declaring that we know the stiffness matrix of the system or, equivalently, the autocorrelation matrix of the underlying random process. Not using any model and just collecting data and using statistics, on the other hand, implies no a priori knowledge of these matrices and instead represents an attempt to approximate them through experimental observation. One important difference, however, is that the FEM and physical modeling provides interpolation functions, while experimental observation does not. This makes physical modeling better suited for resampling and warping, which is important for our application.

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

$$\boldsymbol{\Phi}^T \mathbf{S} \boldsymbol{\Phi} \approx \boldsymbol{\Phi}^T \mathbf{K}^{-1} \mathbf{K}^{-1} \boldsymbol{\Phi} \quad (20)$$

$$= (\boldsymbol{\Phi}^T \mathbf{K}^{-1} \boldsymbol{\Phi}) (\boldsymbol{\Phi}^T \mathbf{K}^{-1} \boldsymbol{\Phi}) \quad (21)$$

$$= \mathbf{A}^{-1} \mathbf{A}^{-1} \quad (22)$$

$$= \begin{bmatrix} \frac{1}{\lambda_1^2} & & & \\ & \frac{1}{\lambda_2^2} & & \\ & & \ddots & \\ & & & \frac{1}{\lambda_{3N}^2} \end{bmatrix} \quad (23)$$

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 $\boldsymbol{\Phi}$ also diagonalizes \mathbf{S} , which implies that $\boldsymbol{\Phi}$ consists of the eigenvectors of \mathbf{S} . Therefore the eigenvectors of \mathbf{S} converge to those of \mathbf{K} , which says that the physical and experimental modes are the same under the assumption of a uniform distribution of forces \mathbf{R} . In other words, the directions of variability in a set of data found by diagonalizing the sample covariance matrix converge toward the free vibration modes of the system. Therefore, for example, the direction of maximum variability in the data is the vibrational mode with the lowest frequency. This makes sense because the low frequency vibrational modes have the largest displacement amplitudes, and therefore over a range of sample objects the low frequency amplitudes will show the most variance.

Third, because of the reciprocal relationship between the eigenvalues of \mathbf{S} and \mathbf{K} , the high-variance directions (large eigenvalues) in a principal component analysis are the low-frequency directions (small eigenvalues) in a modal decomposition.

In the remainder of this section we show how the above analysis allows us to decouple an overall stiffness matrix \mathbf{K} into effects that can be physically modeled and effects that can be experimentally approximated. In Section 4 we will apply this stiffness decoupling to the brain, modeling the low-frequency ventricular shape variation due to overall brain shape, and experimentally approximating the high-frequency ventricular shape variation due to pathology.

3.5 Compound modes

Assume that the deformation we observe in the square is the result of two independent physical processes, so that its total deformation is just a combi-

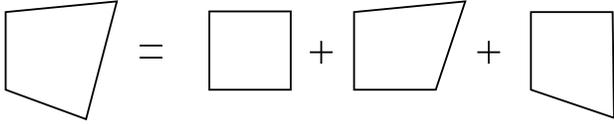


Figure 3: Total deformed shape = original shape + deformation due to known physical process + deformation due to unknown physical process.

nation of two independent deformations, as shown in Figure 3. Because of the equivalence of physical and experimental modes described earlier, we can represent each of these processes as we see fit, without regard to how the other process is being represented. For example if we have good physical models for both processes, then we could use modal analysis for both. Or if we have data available for one, and a good physical model available for the other, then we could use both experimental and physical modes at the same time. This is the situation we will describe here, since later on it will be the one applicable to our brain model.

Assume therefore that we understand the physics of the first process, so that we are able to model it analytically and use the physical modes to describe the deformation that it causes. Also assume that we do not have detailed knowledge of the physics of the second process, so we model its deformation by collecting samples and using experimental modes. If the physics of the first process produces a quadratic energy to be minimized, and if we assume that the second process also has a quadratic energy (but with *unknown* stiffness matrix) then the connection in Section 3.4 applies. Therefore the total deformation is just the sum of the physical deformation modes obtained analytically from the model of the first process and the statistical modes obtained experimentally from the data collected on the second process. Using Equations 12 and 14 we write the total deformation \mathbf{U} as

$$\mathbf{U} = \mathbf{U}^P + \mathbf{U}^E \quad (24)$$

$$= \mathbf{\Phi} \mathbf{a} + \mathbf{\Psi} \mathbf{b} \quad (25)$$

$$= \sum_{i=1}^{3N} \phi_i a_i + \sum_{i=1}^{3N} \psi_i b_i \quad (26)$$

$$= \sum_{i=1}^{6N} \gamma_i c_i \quad (27)$$

$$= \mathbf{\Gamma} \mathbf{c} \quad (28)$$

where

$$\gamma_i = \begin{cases} \phi_i, & i \leq 3N \\ \psi_i, & i > 3N \end{cases}$$

and

$$c_i = \begin{cases} a_i, & i \leq 3N \\ b_i, & i > 3N \end{cases}$$

$\mathbf{\Gamma} = [\phi_1, \phi_2, \dots, \phi_{3N}, \psi_1, \psi_2, \dots, \psi_{3N}]$ is a $3N \times 6N$ matrix whose columns γ_i we will call the *compound* modes of the system.

4 Application of compound modes to brain structures

In this section the compound mode formulation is applied to structures within the brain. We show how the compound modes can be applied in a serial manner in order to improve disease classification. Our representative brain structure will be the ventricular system, although the formulation is also valid for other structures inside the brain, as will be discussed in Section 8.

4.1 Deformation model

We want to be able to classify the ventricle shapes as either normal, or as belonging to one of several known disease classes. Each disease affects the shape of the ventricles in a unique way, and it is this shape difference that we want to use to do the classification. Unfortunately ventricular shape is also affected by overall head shape, which has nothing to do with the disease processes. This is illustrated in Figure 4, which shows schematic drawings of the intra-cranial cavity (ICC) and the ventricles, in various configurations.³ The basic idea is that to a first order, the shape of the ventricles is affected by the shape of the head (justification of this is given in the next section). Therefore, for example, two people with differently shaped heads will tend to have differently shaped ventricles, even in the absence of pathology. This is illustrated in Figure 4a. 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 4b. 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 4c shows this case.

Therefore there is a normal, nonpathological deformation of the ventricles that will interfere when we try to do our classification. Hence we seek to remove the

³The ICC is used instead of the brain because we need a structure that can help in modeling *nonpathological* deformation. Whereas the brain surface may change because of the presence of disease, unless the skull is affected by the disease, the ICC will not change.

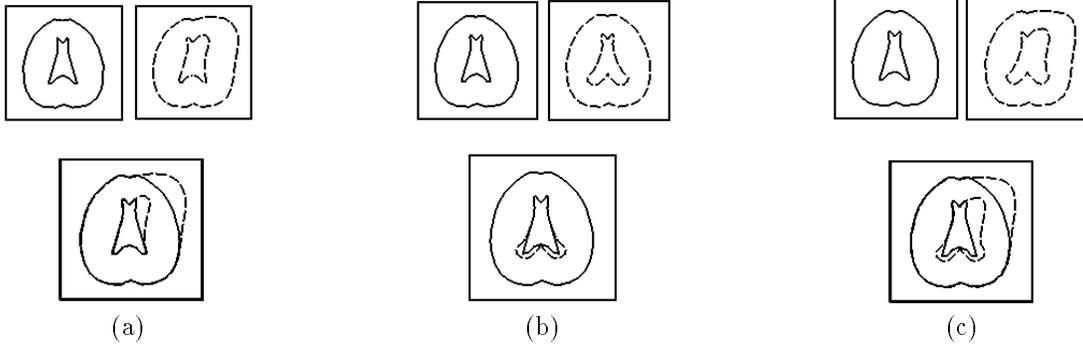


Figure 4: Schematic representation of ICC and ventricles. (a) Different shaped ICCs, no ventricular disease present. The only ventricular shape difference is due to the ICC shape difference. (b) Same shaped ICCs, with ventricular disease. The lower tips of the ventricles are expanded due to the disease. This shape difference is not caused by different ICC shapes. (c) Different shaped ICCs, with ventricular disease. The pathological difference in ventricular shape is partially masked out by the nonpathological difference due to ICC shape.

normal deformation, while leaving in the pathological deformation, and then to do the classification.

Mathematically we can write

$$\mathbf{u}(x, y, z) = \mathbf{u}^G(x, y, z) + \mathbf{u}^L(x, y, z) \quad (29)$$

where $\mathbf{u}(x, y, z)$ is a 3×1 vector of the total deformation at spatial location (x, y, z) , $\mathbf{u}^G(x, y, z)$ is the deformation due to *global* effects (related to the shape of the ICC), and $\mathbf{u}^L(x, y, z)$ is the deformation due to *local* effects (caused by pathology and individual local variation).

Using Equation 29, we can represent a patient’s ventricular system by a list of deformations away from an average ventricular system. This list will contain one entry for every spatial location in the ventricular system that we wish to use in our model. This could be every voxel in the ventricular volume, just the ventricular surface voxels, or even just a small set of ventricular landmarks. The particular choice is an implementation issue. For now we just refer to this set of spatial locations as the *ventricle voxels*. In vector form we have

$$\mathbf{U} = \mathbf{U}^G + \mathbf{U}^L \quad (30)$$

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

$$\mathbf{U} = \begin{bmatrix} \mathbf{u}(x_1, y_1, z_1) \\ \mathbf{u}(x_2, y_2, z_2) \\ \vdots \\ \vdots \\ \mathbf{u}(x_V, y_V, z_V) \end{bmatrix} \quad (31)$$

and \mathbf{U}^G and \mathbf{U}^L are defined similarly. V is the number of ventricle voxels.

4.2 Use physical modes for ICC deformation

4.2.1 Justification: brain growth and ventricular shape

One reason for using physical instead of experimental modes for the global deformation is that we can formulate an approximate physical model for the brain as a whole, and this represents a priori information that should be used. The justification of this comes from developmental neuroanatomy. As the cerebral hemispheres expand during brain growth, the structures within the brain are affected by resulting pressures and forces and hence, to a first approximation, expand in the same way as the hemispheres. This is why, for example, the lateral ventricles are C-shaped. Both the cerebral hemispheres and the lateral ventricles start out approximately spherical. As the hemispheres grow, they project outward and around the brain stem (think of the top of a mushroom expanding to cover the stalk). This approximately C-shaped expansion is transferred to the ventricles through various pushing and pulling forces, resulting in a ventricular shape correlated to the cerebrum shape.

A second and more practical reason is the availability of analytic interpolation functions that allow us to relate deformations at one point to forces and deformations throughout the body. 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.

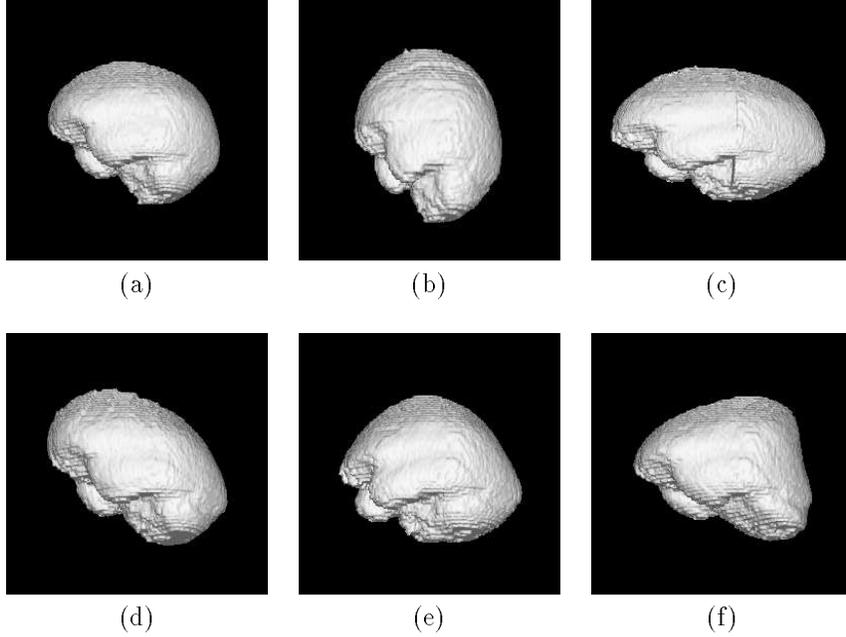


Figure 5: The average ICC (a) and several of the ICC’s physical deformation modes (b-f).

4.2.2 ICC physical modes

Although the cerebrum does exert forces on the ventricles and other structures as it grows, detailed knowledge of this process does not exist. We therefore model the brain as a homogeneous piece of elastic material. This model is discretized by assuming the mass to be concentrated at M nodes. A stiffness matrix \mathbf{K} is then formed using the assigned material properties and the geometry of the nodes. The nodes are chosen in such a way that they are independent of any pathological effects, thus keeping the global and local effects independent of one another.

Figure 5 shows several of the physical deformation modes of the ICC, computed using the “idealized modes” technique described in [2]. 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 [11] will recover modes that are closer to the actual ICC modes.

4.2.3 Application of ICC physical modes to ventricles

By warping one ICC to match the other, the set of M nodal displacements can be found. After the solution is obtained, the displacement at any spatial location

(x, y, z) inside the ICC can be written in terms of the M ICC nodal displacements as

$$\mathbf{u}^G(x, y, z) = \mathbf{H}(x, y, z)\mathbf{U}^{ICC} \quad (32)$$

where \mathbf{H} is the $3 \times 3M$ FEM interpolation polynomial and \mathbf{U}^{ICC} is the $3M \times 1$ vector of the ICC nodal displacements. Using the modal transformation described in Section 3.2, the $3M$ nodal displacements can also be written as a linear combination of the free vibration modes of the ICC:

$$\mathbf{U}^{ICC} = \mathbf{\Phi}\mathbf{a} = \sum_{i=1}^{3M} \mathbf{\Phi}_i a_i \quad (33)$$

where $\mathbf{\Phi}$ is the matrix whose columns consist of the eigenvectors of \mathbf{K} , and the a_i are the associated modal amplitudes.

Plugging Equation 33 into Equation 32 gives

$$\mathbf{u}^G(x, y, z) = \mathbf{H}(x, y, z)\mathbf{\Phi}\mathbf{a}. \quad (34)$$

Applying this equation to each of the V ventricle voxels yields $3V$ equations, which can be combined in vector form as

$$\mathbf{U}^G = \mathbf{H}\mathbf{\Phi}\mathbf{a} \quad (35)$$

where \mathbf{H} is now a $3V \times 3M$ matrix given by

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}(x_1, y_1, z_1) \\ \mathbf{H}(x_2, y_2, z_2) \\ \vdots \\ \mathbf{H}(x_V, y_V, z_V) \end{bmatrix} \quad (36)$$

This is the representation of \mathbf{U}^G in Equation 30 in terms of the physical modes of the ICC.

4.3 Use experimental modes for disease deformation

We were able to use physical modes for \mathbf{U}^G because we had at least an approximate physical model of how the brain as a whole deforms under forces. We have no such luck with the deformation of a particular brain structure due to a particular disease. Consider, for example, the ventricular system, and how it is affected by Alzheimer’s disease. Alzheimer’s disease causes different parts of the brain to atrophy. This decrease in brain matter is partially taken up by an expansion of the ventricles. To create any type of realistic physical model of the ventricular deformation due to Alzheimer’s disease, then, one would have to know which of the surrounding brain structures were atrophying, the geometric relationship of the ventricles to these surrounding structures for the patient at hand, and also to how model the cerebrospinal fluid (CSF) hydrodynamics. Also, there are other diseases (e.g. hydrocephalus) that affect the CSF pressures in ways that are not well understood. All of these unknown factors make physically modeling the pathological ventricular deformations extremely difficult.

Therefore we turn to the experimental modes. First we note that, by definition (see Equation 29), $\mathbf{u}^L(x, y, z)$ does not contain any deformation due to head shape. Therefore the starting point when analyzing the experimental modes will be the ventricle voxels, already warped according to Equation 35.

The deformation of the V ventricle voxels away from their average positions is represented by \mathbf{U}^L . Assuming we have P data sets available, we form the sample covariance matrix \mathbf{S} of these deformations (Equation 13). As described in Section 3.3, diagonalizing \mathbf{S} gives the $3V$ experimental modes of variation, which we assemble as the columns in a matrix $\mathbf{\Psi}$. Any ventricle data set, originally represented as a vector of V displacements (each of which has been warped by the physical modes), can instead be written as a linear

combination of these experimental modes:

$$\mathbf{U}^L = \mathbf{\Psi} \mathbf{b} = \sum_{i=1}^{3V} \mathbf{\Psi}_i b_i. \quad (37)$$

4.4 Compound modes

Plugging Equations 35 and 37 into Equation 30 then gives a decomposition of our original ventricle data set into $3M + 3V$ *compound* modes:

$$\mathbf{U} = \mathbf{U}^G + \mathbf{U}^L \quad (38)$$

$$= \mathbf{H} \mathbf{\Phi} \mathbf{a} + \mathbf{\Psi} \mathbf{b} \quad (39)$$

$$= \mathbf{H} \sum_{i=1}^{3M} \mathbf{\Phi}_i a_i + \sum_{i=1}^{3V} \mathbf{\Psi}_i b_i \quad (40)$$

$$= \sum_{i=1}^{3M} \mathbf{H} \mathbf{\Phi}_i a_i + \sum_{i=1}^{3V} \mathbf{\Psi}_i b_i \quad (41)$$

$$= \sum_{i=1}^{3M+3V} \gamma_i c_i \quad (42)$$

$$= \mathbf{\Gamma} \mathbf{c} \quad (43)$$

where $\mathbf{\Gamma}$ and \mathbf{c} are just the three-dimensional analogs of the $\mathbf{\Gamma}$ and \mathbf{c} defined in Equation 28, with the physical modes now interpolated through the use of \mathbf{H} .

At this point it is instructive to examine what exactly Equation 38 allows us to do. If we do not decompose \mathbf{U} into its ICC and pathological components, then we would have a total of just $3V$ modes:

$$\mathbf{U} = \sum_{i=1}^{3V} \mathbf{\Theta}_i d_i. \quad (44)$$

where the $\mathbf{\Theta}_i$ and d_i are different than the $\mathbf{\Psi}_i$ and b_i of Equation 40 because no warping takes place here before calculating the modes.

Suppose, for example, that only two modes were seen in the data - an overall expansion of the ventricles, and a skewing of the ventricles. Furthermore, assume that we are analyzing two diseases - disease #1 which causes a pure expansion of the ventricles, and disease #2 which results in a pure skewing of the ventricles. Also assume that, unfortunately for disease discrimination, head shape can also be found to come in varieties that are described by a certain amount of expansion/contraction, and a certain amount of skewness. Assuming that this affects ventricular shape, then each patient’s ventricles will also have an effect due to this, independent of any disease which may be present. Therefore taking a collection of patients, calculating the experimental modes of the ventricles,

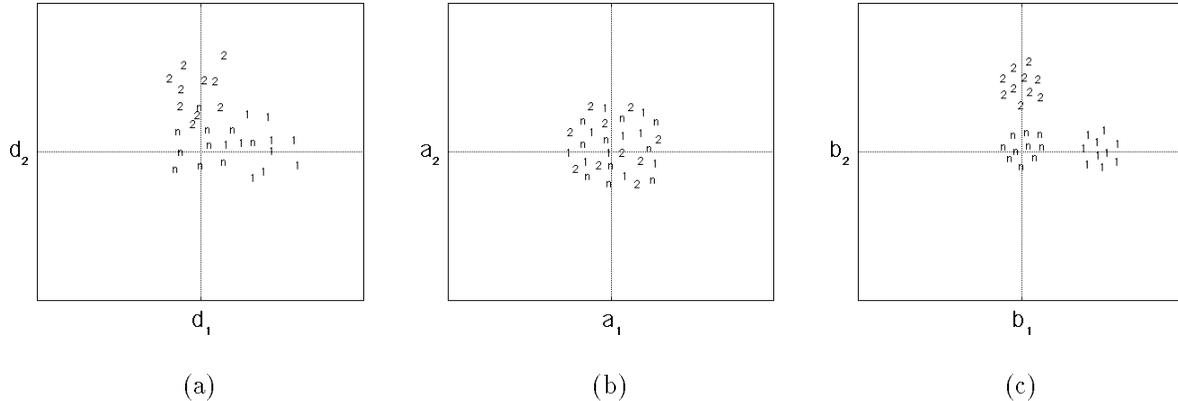


Figure 6: Projection of hypothetical patients onto modes. A 1 refers to a patient with disease #1, a 2 refers to a patient with disease #2, and an “n” refers to a normal, healthy person. (a) Original experimental modes. (b) Physical modes of ICC. (c) Physically-discounted experimental modes.

and then projecting each patient onto these two modes would look something like Figure 6a, where we have constructed a hypothetical set of patients. A 1 refers to a patient with disease #1, a 2 refers to a patient with disease #2, and an “n” refers to a normal, healthy person. We see that the presence of differing head shape makes disease discrimination difficult.

The procedure just described corresponds to Equation 44, which is why the axes are labeled d_1 and d_2 . Equation 38 allows us to improve upon this. Suppose that we are able to account for head shape *before* calculating the experimental modes. Hence we now have broken up the ventricular displacements into four directions of variability instead of just two - ventricle size due to head size, ventricle size due to disease, ventricle skewness due to head skewness, and ventricle skewness due to disease. Since the diseases do not affect head shape, then the subspace spanned by the two head shape directions is orthogonal to the subspace spanned by the two disease directions.

We now have four orthogonal directions to work with instead of two. Equation 38 tells us to disregard the two head shape directions, and just use the two disease directions to perform the discrimination. Figure 6b shows the projections onto the two head shape directions. Since these directions are uncorrelated with the diseases, we have just a random distribution for the diseases and for the normals. Figure 6c, however, shows the improved class separability once the ICC modes have been discounted. The labels of the axes in Figures 6b (a_1 , a_2) and 6c (b_1 , b_2) are the mode coefficients of Equation 40.

Note that Figure 6 is just another way of drawing Figure 3 for many patients at once - i.e. the total

deformation of a patient’s ventricles (Figure 6a) is just the sum of the random ICC effects (Figure 6b) and the pathological effects (Figure 6c).

Ventricles have a complicated shape, and the way they deform is also complicated. Medically, however, shape differences due to diseases are known to exist. Therefore, in some high-dimensional space there exist directions that will describe the variation caused by each disease. Equation 38 explicitly tells you what those directions are. The whole trick, of course, is in finding the components on the right hand side of Equation 38. Unfortunately the real world is not like our toy example. There will be many more modes of variability for both head shape and disease effects, the effect that a disease has on the ventricles will not be completely unique to that disease, and there will be other sources of noise in the data and in the processing. In Section 6 we will see the effects of all of this.

4.5 Example: Healthy subject with large cranium

Figures 7 and 8 show an example where the warping of the cranium contents results in ventricles that more closely resemble average ventricles. Figures 7a and 7b show the average ICC and ventricles, while Figures 7c and 7d show the ICC and ventricles of one of the healthy subjects (the ventricles are scaled up in size relative to the ICCs). This 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 Figures 7b and 7d. Calculating the ICC physical deformation modes that warp Figure 7c into Figure 7a, and then applying

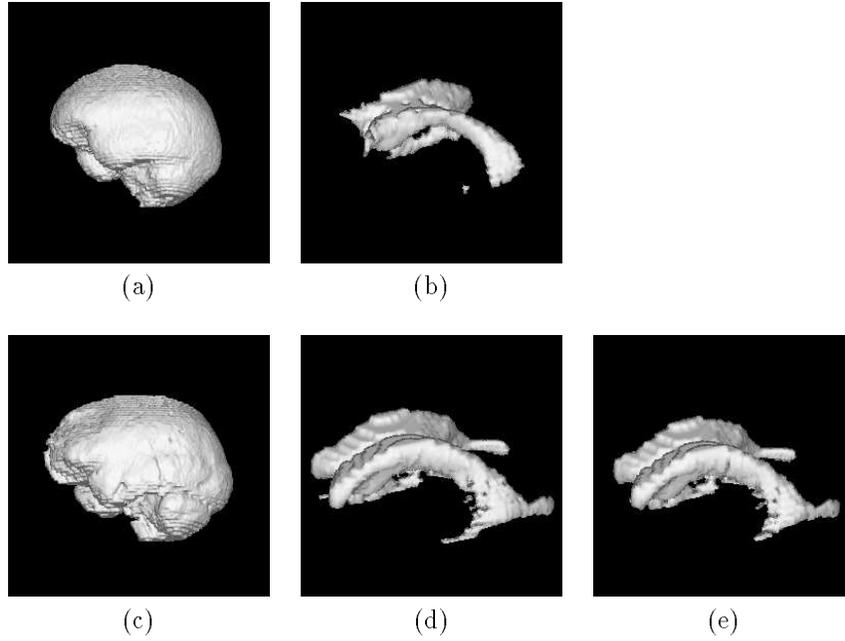


Figure 7: Normalizing ventricular shape by cranium shape. In both this and the next figure, the size of all the ventricles has been increased for clarity. However all the ICCs are scaled consistently with one another, and all the ventricles are scaled consistently with one another. (a) Average ICC (computed 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 his 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).

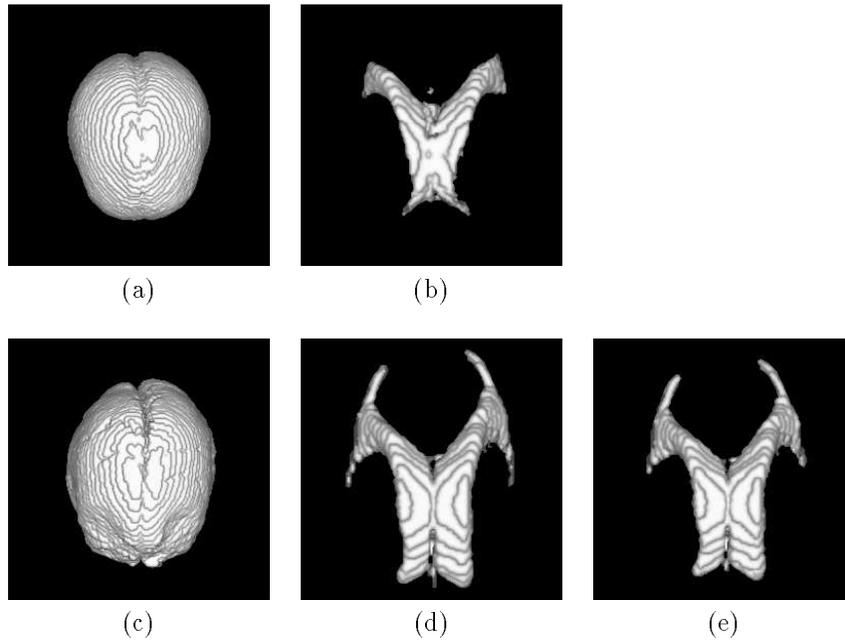


Figure 8: Same structures as in the above figure, but viewed from above. The front-to-back direction is now bottom-to-top. See the caption of the above figure for details.

that warping to the ventricles in Figure 7d produces the warped ventricles in Figure 7e. These warped ventricles have become more similar to the average ventricles in Figure 7b, because we have taken into account the different head shapes.

5 Modes of variation

Figure 9 illustrates some of the ventricular modes of variation obtained from the 25 patients. Figure 9a shows the average ventricular system obtained from the healthy subject, while the first mode of variation among the normals is shown in Figure 9b. This mode shows that overall ventricular volume and length, as well as the size of the posterior horns, are important features in discriminating among normal ventricles.

Figures 9c and 9d are the average AD ventricular system and the first AD mode of variation. Some enlargement in the average AD is seen relative to the average normal. Similar to the normals, the first AD mode indicates that the size of the posterior horns is important. The size of the inferior horns also appears to be significant.

The average NPH ventricular system and the first mode of variation among the NPH ventricles are shown in Figures 9e and 9f. The most important features are the large expansion and an overall puffy shape.

6 Experiments and results

Each of the twenty-five patients underwent an MR brain scan. These scans were then postprocessed to first segment out the relevant parts (the ventricles and the entire ICC), and then the surface voxels of these parts were extracted. After a rigid body alignment of the 25 ICCs, our method was applied. First, the principal components of the ventricle surface voxel displacements were calculated. Second, for each patient the physical deformation modes that warp the patient’s ICC to the average ICC were applied to the patient’s ventricles. The principal components of these *warped* ventricle surface voxel displacements were then computed. This provides us with the experimental modes of ventricular shape deformation, both with (as in Figure 6c) and without (as in Figure 6a) first discounting them with the physical modes of the ICC.

The experimental modes of the ventricles form a very high-dimensional feature space, since originally there is one dimension for every ventricle surface voxel. However, as there are only 25 patients, they will live in a 24-dimensional subspace of the original feature space. Performing the eigenanalysis produces 24 vec-

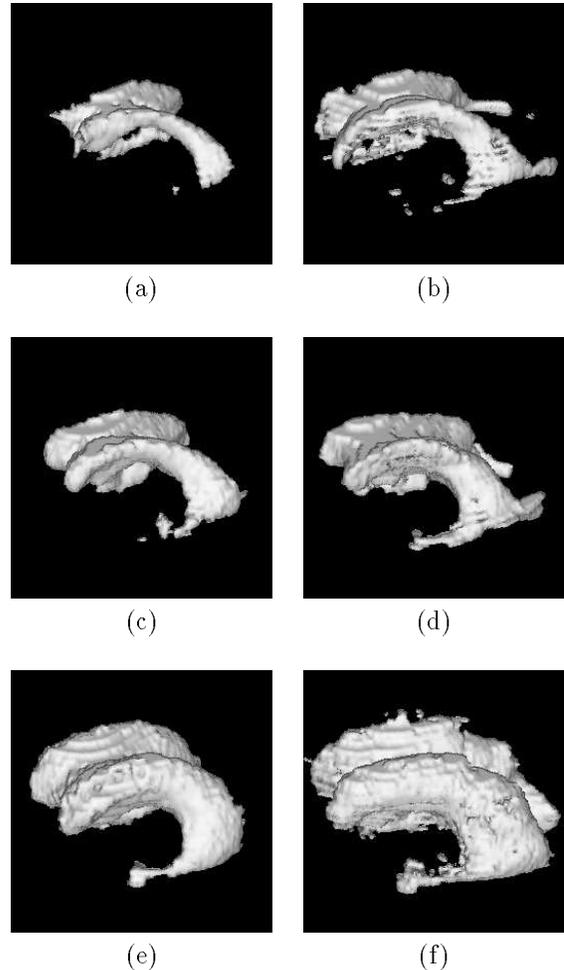


Figure 9: (a) The average normal ventricular system. (b) The first normal mode of variation. (c) The average Alzheimer’s disease ventricular system. (d) The first mode of variation caused by Alzheimer’s disease. (e) The average normal-pressure hydrocephalus ventricular system. (f) The first mode of variation caused by normal-pressure hydrocephalus.

tors (the 24 eigenvectors of the covariance matrix that correspond to nonzero eigenvalues) that span this subspace. Noise considerations imply that the higher-order eigenvectors are not reliably estimated; an engineering rule of thumbs states that only the lower $\frac{1}{4}$ eigenvectors are well estimated.⁴ Therefore we keep only the directions given by the first six eigenvectors. These six eigendirections span a six-dimensional space that represents measurements on our 25 patients that

⁴For instance if we took another set of 25 patients and computed their 24 eigenvectors, the first six eigenvectors would be approximately the same as the six we computed here, whereas the higher-order ones would be different.

Feature(s)	# Correct
Eigenanalysis	24
Volume	22

Table 2: Number of patients correctly classified (out of 25), using the first 6 experimental modes of the unwarped ventricles as features. Also shown are the classification results obtained using just ventricular volume.

show large variation from patient to patient.

6.1 Physical modes not discounted

In the introduction we said that our method can be considered to be a more sophisticated version of using just ventricular volume, or ventricular volume corrected for by ICC volume. By calculating the experimental modes without first subtracting off the ICC’s physical modes, we are just performing an eigenanalysis on the original ventricle data. If ventricular volume helps discriminate between patients, then it will be represented in the calculated eigenvectors.

Table 2 shows the result of applying the quadratic classifier to the experimental modes of the unwarped ventricle data. For comparison, the results obtained using just ventricular volume, shown earlier in Table 1, have been appended to this table. Some improvement in performance is suggested when using the experimental modes.

6.2 Physical modes discounted

We now examine the application of both phases of our model to the 25 data sets. By warping all 25 ICCs to the same reference shape, and then applying this warping to the ventricles, some of the ventricular shape variation due to head shape should be removed. Performing the eigenanalysis on these warped ventricles should therefore result in a six-dimensional space in which the classes have better separability than if the warping is not performed.

6.2.1 Mean square distance

One measure of success is the mean square distance between the average ventricular system and each of the patient’s ventricles. This is measured by computing the distance between each voxel on the average ventricular surface and the nearest point on each patient’s ventricular surface. The sum of these for a

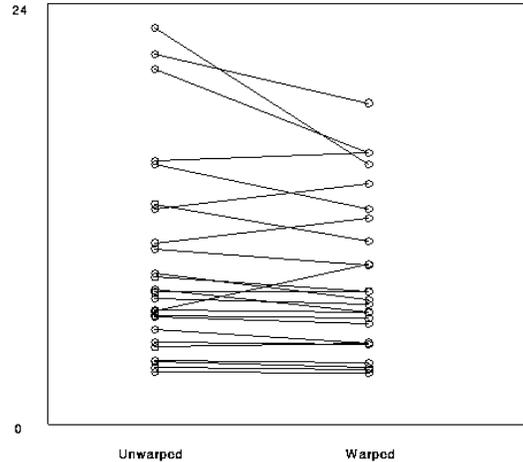


Figure 10: Mean square distance (in mm^2) between ventricles before and after warping

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 our results. In 19 of the 25 patients, the mean square distance decreased with warping. A paired t -test gives a significance level of 0.06 for rejecting the hypothesis that the two means are equal. This indicates that the ICC warping is making most patients’ ventricular systems more like the average - i.e. it is removing some of the ventricular shape variation due to differing head shapes.

Feature(s)	# Correct
Normalized eigenanalysis	25
Eigenanalysis	24
Volume	22
Normalized volume	20

Table 3: Number of patients correctly classified (out of 25), using the first 6 experimental modes, both with and without first discounting them by the physical modes of the ICC. Also shown are the classification results obtained using just ventricular volume.

6.2.2 Classification

Table 3 shows the results of running the Gaussian classifier on the experimental modes of the warped ventricle data. Comparing the warped to the unwarped

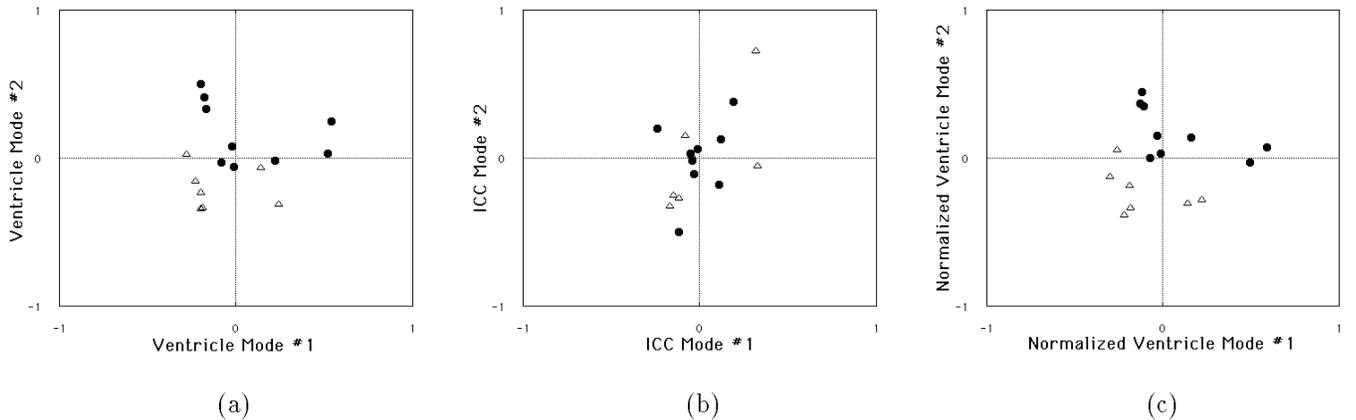


Figure 11: Projection onto modes computed from normals and NPH patients. Circles are healthy subjects and triangles are patients with NPH. (a) Ventricular deformation modes, computed without normalizing ventricular shape by cranium shape. (b) ICC physical modes. (c) Ventricular deformation modes, computed after first discounting by the physical modes of the ICC.

eigenanalysis, we see that the number of patients correctly classified improves from 24 to the full 25. To test of the robustness of this result, we varied the number of eigenvectors used in the classification, as well as the type of classifier. In almost all of the cases the best results were obtained when using the physically-discounted experimental modes as features.

We have also appended the results of Table 1. As we stated earlier, discounting by the physical modes means that we are normalizing the ventricular shape by the overall head shape, and thereby removing some nonpathological ventricular shape variance. Whereas dividing ventricular volume by ICC volume takes into account just head size, we are taking into account the complete head shape. Comparing the results, we see that the use of the warped eigenvectors results in a feature space in which it is easier to draw decision boundaries that separate the classes.

Since the results in Tables 2 and 3 represent effects taking place in six dimensions, they are hard to visualize. We therefore took just the normal subjects and the NPH patients, and performed an eigenanalysis both with and without first warping based on the ICCs. The projections onto just the top two eigenvectors are shown in Figures 11a and 11c, while the projections onto the top two ICC deformation modes are shown in Figure 11b. The projections onto the ventricular deformation modes (Figures 11a and 11c) show some separation between the two classes, while the projections onto the ICC modes are distributed randomly. This is to be expected, since cranium shape is uncorrelated with disease state. Note that Figures

11a-c are the real-data analogs of Figures 6a-c. Although not as pronounced as the fake example of Figure 6, an improvement in class separability can be seen when discounting by the physical modes of the ICC (Figure 11c vs 11a).⁵

7 Open issues

7.1 Model accuracy

One area requiring further study is the accuracy of physical models of the ICC. Flaws in these models will cause errors in the warping of the cranial contents, resulting in suboptimal shape normalization for the ventricles and other interior structures.

The creation of these models has long been an active research area among biomechanical investigators, who have constructed both solid and fluid mechanic models [12][13][14][15]. The constitutive parameter values plugged into these models have come from experiments on human cadavers and both living and dead animals. Due to the very complex nature of brain tissues and difficulties in extrapolating cadaver and animal results to living humans, however, the material properties of the brain are still very poorly understood [13][15]. Because most of these models were created

⁵Out of all the possible combinations of using either two or all three classes, we chose to show the combination whose class separability showed the greatest improvement when warping the ICCs. This will cause Figure 11 to *overstate* the effect. However we are looking at only a two-dimensional subspace of the six-dimensional space in which the effect is taking place. This will cause Figure 11 to *understate* the effect.

in order to study head injuries, model validation has typically been performed by running 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 [16][17][4], typically with the goal of performing some type of inter-patient warping. The mechanical properties used for these models have usually been set to those of much simpler materials. In our model, for instance, we build the stiffness matrix using extremely elastic material properties. It is doubtful, however, that much improvement in stiffness matrix accuracy would be obtained by employing the material property estimates found in the biomechanics literature, as these are highly uncertain.

By applying the principles outlined in this paper, we can forgo the reliance on measured material properties 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. Because the ICC of each patient is sampled at different locations by the medical imaging scanner, this ICC displacement data will be irregularly sampled. Fortunately, the FEM’s physical interpolation functions can be used to refer irregularly-sampled displacement data back to the FEM nodes of the model. Once at these nodes, Equations 15-19 can then be used to estimate the stiffness matrix \mathbf{K} .

Here we see the real power of the FEM and Equations 15-19. The FEM’s physical interpolation functions allow you to convert irregularly sampled ICC displacements into regularly-sampled (at the FEM nodes) displacement data in a physically meaningful fashion. Then Equations 15-19 tell you how to convert this regularly-sampled displacement data into a stiffness matrix.

Hence we have two ways to construct the stiffness matrix. The one used in this paper relies on experimental observation of material properties, and then builds the stiffness matrix in the usual FEM fashion. The method proposed in this section creates the stiffness matrix directly from experimental observations of ICCs, via Equation 19. Which method creates a more reliable stiffness matrix depends on the accuracy of the

measured material properties of the ICC, the sample size used in estimating the stiffness matrix directly, and the validity of the assumptions used in deriving Equation 19.

Whether the first or second method is used, or even if the two sources of information are combined, the goal is to produce the most physically-correct stiffness matrix. This will provide the most realistic ICC deformation modes, and will give the best shape normalization when the cranial contents are warped.

7.2 Ventricular/ICC shape correlation

Because the shapes of the ICC and interior structures are not perfectly correlated, even a totally accurate physical model of the ICC will not produce complete shape normalization when warping interior structures. While some correlation exists, the exact amount and types of it are not known. In the example of Section 4.5, for instance, shape correlation was present and we were able to successfully remove some of it. There are cases, however, when there is little or no correlation; warping the cranial contents according to ICC shape will therefore not help in normalizing ventricular shape.

Figures 12 and 13 provide an example where head shape normalization does not increase ventricular shape similarity. Even though the subject’s ventricles (12d) are significantly larger than average (12b), his ICC (12c) is not that different from average (12a). Using the ICC shape difference to warp his cranial contents, therefore, does not produce ventricles (12e) that are any closer to the mean.

To discover which structures will benefit most from ICC-based warping, models of the correlation between interior structure shape and ICC shape are needed. Therefore in addition to obtaining just the ICCs of a large normal population, as suggested earlier, the shapes of interior structures should also be extracted. The shape correlations between these structures and the ICCs can then be calculated and used to decide which structures will benefit most from warping. In addition, it may be possible to construct warping functions that take into account the particular types of shape correlations found.

7.3 Correspondence

Computing the various deformation modes requires knowing the correspondence between points on surfaces. In this paper we used a closest-match technique to approximate this correspondence. For calculating the correspondence between two ICCs, this procedure is adequate. For structures as complicated as

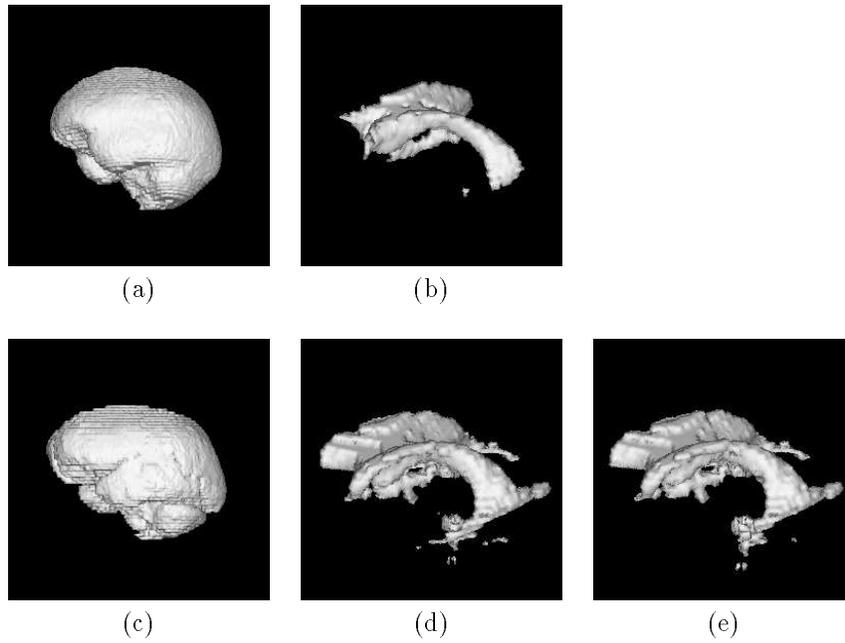


Figure 12: An example where normalizing ventricular shape by cranium shape does not improve matters. (a),(b) Average ICC and normal ventricles. (c),(d) ICC and ventricles of one of the healthy subjects. The ventricles are far from average, but the ICC is close to average. Therefore warping this subject's ventricles using ICC shape differences does not produce ventricles (e) any closer to average.

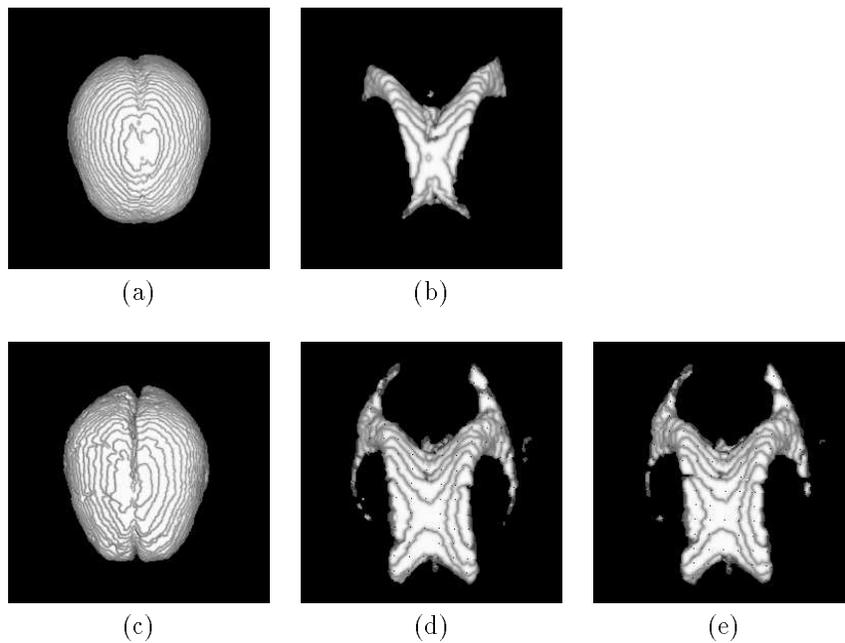


Figure 13: Same structures as in the above figure, but viewed from above. The front-to-back direction is now bottom-to-top. See the caption of the above figure for details.

the ventricular system, however, closest-match techniques will provide only a very coarse approximation to the true correspondence. Any errors in this correspondence will propagate into inaccuracies into the estimated ventricular deformation modes.

Possible solutions to this problem include the use of more sophisticated point correspondence algorithms, or the extraction of more salient features from the ventricles. This feature extraction could be done manually, having a user simply pick corresponding landmarks on the different ventricular data sets. Aside from the unattractiveness of introducing user intervention, we have found that it is very difficult even for medical experts to reliably pick homologous points on different realizations of a complex structure. A more promising approach is to automatically detect these points or lines using concepts from differential geometry [18].

8 Summary

We have presented a new method that describes shape deformation by using both physical modeling and statistical techniques. When applied to the human brain, our method is able to separate out pathological shape deformation 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 experimental modes of a brain structure by the physical modes of the intracranial cavity, it may be possible to improve disease classification.

Although the results presented were only for the ventricular system, our technique is general and therefore applicable to other structures within the brain. Any structure that is affected by the growth of the cerebral hemispheres during neuroanatomical development, in such a way that its shape is positively correlated in some manner with overall brain shape, can be modeled using our method.

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 morphology is correlated to overall brain shape. The first two of these can be improved upon with more data and better implementation; overcoming the last limitation requires a model of whatever process is thought to be affecting the morphology.

In summary, there are three main contributions of

this work. First, we have suggested a shape representation method for brain structures that we believe to be intuitive, in the sense that it can be thought of as a more sophisticated version of the simpler idea of using volume and head size to describe and normalize the shape of a brain structure. Second, by using the connection between physically-based modeling and probability, we have put our method into a consistent mathematical framework. This enabled us to decouple the modes of variation caused by different physical processes, and to use these decoupled modes in standard pattern recognition algorithms. Third, by applying our method to medical data, we have indicated that it may be possible to improve disease classification by discounting the experimental modes with the physical modes associated with head shape. These experimental modes were displayed as an illustration of the ventricular deformation modes caused by Alzheimer's disease and by normal-pressure hydrocephalus.

Acknowledgments

The authors would like to thank Stan Sclaroff and Irfan Essa for their helpful advice, and Mike Sasmuae for providing the MR brain data.

References

- [1] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models and 3D object reconstruction. *Int. J. Comput. Vision*, 1:211–221, 1987.
- [2] 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.
- [3] F. L. Bookstein. *Morphometric Tools for Landmark Data: Geometry and Biology*. Cambridge University Press, 1991.
- [4] D. L. Collins, T. M. Peters, W. Dai, and A. C. Evans. Model based segmentation of individual brain structures from MRI data. In *SPIE: Visualization in Biomedical Computing*, pages 10–23, 1992.
- [5] M. Turk and A. P. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.
- [6] 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.

- [7] 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.
- [8] K. Bathe. *Finite Element Procedures in Engineering Analysis*. Prentice-Hall, Inc., 1982.
- [9] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intell.*, PAMI-6(6):721–741, November 1984.
- [10] R. Szeliski. *Bayesian Modeling of Uncertainty in Low-level Vision*. Kluwer Academic Publishers, 1989.
- [11] S. Sclaroff and A. P. Pentland. On modal modeling for medical images: Underconstrained shape description and data compression. In *Proc. Workshop on Biomedical Image Analysis*, June 1994.
- [12] M. Chan, C. Ward, D. Schneider, and S. Adams. Relative importance of skull deformation. In *Proc. Biomechanics Symposium*, 1981.
- [13] C. Chu, M. Lin, and M. Lee. Finite element analysis of cerebral contusion. *Journal of Biomechanics*, 27(2):187–194, 1994.
- [14] W. Goldsmith. *Biomechanics, Its Foundations and Objectives*. chapter Biomechanics of Head Injury. Prentice-Hall, Inc. 1972. Y.C. Fung, N. Perrone, and M. Anliker, Ed.
- [15] J. S. Ruan, T. Khalil, and A. I. King. Dynamic response of the human head to impact by three-dimensional finite element analysis. *Journal of Biomechanical Engineering*, 116:44–50, 1994.
- [16] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Comput. Vision Graphics Image Process.*, 46:1–21, 1989.
- [17] C. E. Christensen, M. I. Miller, and M. Vannier. A 3D deformable magnetic resonance textbook based on elasticity. In *Proc. AAAI Symposium on Applications of Computer Vision in Medical Image Processing*, March 1994.
- [18] O. Monga and R. Lengagne. Extraction of the zero-crossing of the curvature derivative in volumic 3D medical images: A multi-scale approach. In *Proc. AAAI Symposium on Applications of Computer Vision in Medical Image Processing*, March 1994.